

Matching Auctions*

Daniel Fershtman[†]

Alessandro Pavan[‡]

November 2017

Abstract

We study mediated many-to-many matching in markets in which valuations evolve over time as the result of shocks, learning through experimentation, or a preference for variety. The analysis uncovers the key trade-offs that platforms face in the design of their matching protocols. It shows that the dynamics that maximize either the platform’s profits or welfare can often be sustained through auctions implementing the matches with the highest bilateral score up to capacity. In equilibrium, bidding is straight-forward and myopic. The analysis also sheds light on the merits of regulating such markets. When match values are positive, profit maximization involves fewer and shorter interactions than welfare maximization. This conclusion need not extend to markets where certain agents dislike certain interactions.

JEL Classification Numbers: D82, C73, L1.

Keywords: matching, experimentation, platforms, bandit problems, asymmetric information, learning, dynamic auctions.

1 Introduction

In the last few years, matching markets have been growing at an unprecedented rate, reflecting the role that the “sharing economy”¹ is taking in the organization of modern business activities. In electronic commerce, for example, a sizeable fraction of trade is mediated by business-to-business (B2B) platforms matching vendors with procurers in search of business opportunities. Likewise, a sizeable fraction of online advertising is mediated by media outlets, online malls, videogame consoles, search engines, and ad exchanges, matching advertisers

*A previous version circulated under the title “Re-Matching, Experimentation, and Cross-Subsidization.” For helpful comments and suggestions, we thank Dirk Bergemann, Yeon-Koo Che, Eddie Dekel, Drew Fudenberg, Benjamin Hermalin, Terrence Johnson, Fuhito Kojima, Larry Samuelson, Rakesh Vohra, Jean Tirole, Asher Wolinsky, Leeat Yariv, and seminar participants at the University of California Berkeley, the University of British Columbia, UCSD, the University of Chicago, Columbia University, Duke University, Harvard University, the University of Illinois at Urbana Champaign, IESE, MIT, the University of North Carolina, the Toulouse School of Economics, the University of Southern California, Stanford University, the University of Western Ontario, the 2015 Paris Workshop on Information Economics and Dynamics, the 2016 Warwick Economic Theory Conference, the 2016 North American Summer Meetings of the Econometric Society, the 2016 NYUAD Conference on Economic Theory, the 2016 World Congress of the Game Theory Society, and the 2017 Workshop on Matching with Preferences (MATCH-UP). Pavan thanks the National Science Foundation for financial support under the grant SES 1530798 and Bocconi University for its hospitality during the 2017-2018 academic year. The usual disclaimer applies.

[†]Northwestern University. Email: dfershtman@u.northwestern.edu.

[‡]Northwestern University. Email: alepavan@northwestern.edu.

¹Also referred to as “access economy,” or “on-demand” economy.

with either consumers or content providers (see, e.g., “Marketing in the digital age: A brand new game”, *The Economist*, August 29, 2015).

Mediated matching plays an important role also in the growing market for scientific outsourcing. Arrangements in the form of “contract-experiments,” in which researchers or startups contract with laboratories that carry out research on their behalf, have existed for a long time.² Recently, however, new intermediaries, such as Science Exchange, have revolutionized the market by introducing the ideas of the sharing economy. These intermediaries match labs with idle equipment with research units that wish to conduct experiments off-site (see, e.g., “Uber for Experiments,” *The Economist*, December 6, 2014). They introduce parties that otherwise would be unlikely to come in contact, and also provide crucial services that remove most of the complexities involved in such ad-hoc relationships. The participating labs, which are highly differentiated in size, equipment, and specialization, typically allow numerous firms to conduct experiments simultaneously within their facilities. Likewise, firms typically seek to run a battery of experiments across different labs within the same period.³ Over time, labs learn about the reliability of the participating research firms and firms learn about the characteristics of the participating labs. Natural capacity constraints preclude the possibility of all firms using all labs simultaneously. As a result, experimentation and re-matching play an important role in such exchanges. Firms change labs in response to the evolution of their needs and the information they gather about the participating labs. The individual prices asked by the labs to the participating research units are also dynamic, reflecting the information the labs obtain about the participating research units, as well as the competition among the research labs.

Other matching markets sharing similar features include advertising where ad exchanges such as AppNexus, AOL’s Marketplace, and Microsoft Ad Exchange use auctions to dynamically match advertisers with content providers, project finance, where consulting firms match startups with lenders;⁴ lobbying, where commercial firms mediate the interactions between policy makers and interest groups;⁵ the market for private medical-tourism services, where intermediaries match patients from abroad seeking specialized treatments with local physicians providing such treatments;⁶ the market for organized events, where online platforms such as meetings.com match clients in search of conference venues, meeting spaces, corporate hotel travel plans, or other hospitality services, with hotels and venues;⁷ the market for delegated tasks, where platforms such as TaskRab-

²For example, facilities that own wind tunnels perform tests for the aerospace industry, for architectural design, for auto and ship design, and for bicycle design.

³The list of laboratories that work with Science Exchange includes 75 of the top 100 recipients of NIH grant funding. Examples include Johns Hopkins University, the Mayo Clinic, and Harvard Medical School.

⁴A similar role is played by peer-to-peer lending and crowdfunding platforms such as Prosper and LendingClub. Such platforms match borrowers with individual or institutional lenders. In addition to connecting borrowers with lenders, these platforms verify the borrowers’ identities and personal data and manage all stages of funded loans. Prosper operated as an online auction marketplace between 2006 and 2009. It recently switched to a system of pre-set rates determined by an algorithm evaluating borrowers’ credit risk.

⁵See Allard (2008) and Kang and You (2016) for how lobbying firms provide tailored (many-to-many) matching services and dynamically price-discriminate each side of the market. See also Dekel, Jackson, and Wolinsky (2008) for a detailed account of how intermediaries help buying and selling votes.

⁶For example, MEDIGO matches providers of medical services abroad with potential patients. It also offers a “Custom A-to-Z Concierge Package” including a variety of costly services. See the New York Times article <http://www.nytimes.com/2013/08/07/us/the-growing-popularity-of-having-surgery-overseas.html> for details about the growing popularity of overseas surgeries.

⁷As in the case of most other platforms, meetings.com does not simply introduce the parties; it also offers tailored services, such as site selection, contract negotiations, and on-site event management. Matching is many-to-many, as hotels and venues hold

bit match users with “runners” and “handymen” supplying painting, errands, cleaning, and various others temporary services.

All of these markets are intrinsically dynamic, due to the gradual resolution of uncertainty about match values, shocks that alter the desirability of existing matching allocations, or simply a preference for variety. Matching is many-to-many and mediated by profit-maximizing platforms collecting payments from the various sides of the market. Agents change partners using the same platform, and prices vary over time in response to the evolution of the agents’ preferences and information. In most cases of interest, the services provided by the matching intermediary extend well beyond simply introducing agents from the various sides of the market. For example, Science Exchange creates an agreement governing the nature of the interaction between the parties, a service without which negotiations, especially for startups, would be costly and lengthy. It also offers the services of a dedicated scientist with a master’s degree or a PhD, who follows up the entire matching process. The platforms’ costs of such auxiliary services are also naturally dynamic and match-specific, and capacity constraints may restrict the number of interactions that may be accommodated.⁸

How should matching allocations and pricing in such markets respond to the dynamics of market conditions, variations in preferences, and/or the arrival of information? How do matching services provided by private (profit-maximizing) intermediaries compare to those provided by public (welfare-maximizing) intermediaries? Finally, how do the answers to the above questions depend on different market conditions? These questions are receiving a lot of attention from market designers and policy makers.

In large markets such as those for ride services, intermediaries such as Uber and Lyft price discriminate as a function of variations in aggregate demand and supply. In markets with a smaller number of participating agents such as those for on-line advertisement, or scientific outsourcing, analysts have predicted that intermediaries will eventually resort to auctions to dynamically allocate their limited capacity. In fact, auctions are already used in the market for personalized display ads by online search engines such as Google and Yahoo!. These intermediaries have been using variations of the second-price auction, the so-called GSP auction (Generalized Second Price), to match ads with viewers and/or content providers — see, e.g., Edelman, Ostrovsky, and Schwarz (2007), and Gomes and Sweeney (2014). Similarly, auctions are used by online ad exchanges such as Google’s DoubleClick and Microsoft’s Exchange to match advertisers with content providers (see, e.g., Mansour, Muthukrishnan, and Nisan (2012)). Such auctions, however, have been criticized for being static, and in particular for not taking into account the value that the platform, the advertisers, and the content providers assign to learning the match values through experimentation (see, e.g., Li, Mahdian, and McAfee (2010) and the discussion therein).

In this paper, we investigate the key trade-offs that platforms face in the design of their matching protocols. We introduce and then study a class of dynamic matching mechanisms that account explicitly for the fact that, in many environments of interest, match values evolve over time, either exogenously (as the result of changes in preferences and needs), or endogenously (as the result of learning and experimentation). We investigate

multiple events within the same period, and clients (which include corporations and government institutions) seek a variety of locations for different types of events. Matching is also dynamic as the clients’ preferences as well as those of the venues typically change over time as the result of previous experiences and variations in needs and availability.

⁸Medical-tourism intermediaries are limited by the size of the medical facilities they contract with. Lobbying firms are limited in the time they can devote to follow up on individual relationships. Online content providers cannot place too many ads on the same page.

the matching dynamics under profit-maximization and compare them with their counterparts under welfare maximization. Dynamics in our model stem from changes in actual or perceived match values rather than the arrival and departure of agents to and from the market. To the best of our knowledge, this is the first paper to study matching dynamics in environments in which agents revise their beliefs about match values and change partners multiple times.

The key ingredients of our model are the following. The payoff that each agent derives from being matched to any other agent from the opposite side is governed by two components: a time-invariant vertical characteristic that is responsible for the overall importance the agent assigns to interacting with agents from the opposite side of the market; and a vector of time-varying relation-specific values capturing the evolution of the agent’s information and preferences for interacting with specific partners. The latter values evolve stochastically over time and may turn negative, reflecting the idea that certain agents may dislike certain interactions. We consider both the case in which the relation-specific values evolve exogenously, as well as the case in which they evolve endogenously as a function of previous interactions. The latter case may capture either the possibility that agents gradually learn the attractiveness of their partners via individual interactions, or a preference for variety, by which agents’ values change (possibly stochastically) with the number of past interactions. Both the vertical and the horizontal components are the agents’ private information. The agents learn the vertical types prior to entering into negotiations with the platform and learn the horizontal types over time either exogenously or after interacting with other agents. The decomposition of the match values into vertical and horizontal types permits us to shed light on the distortions brought in by the platforms’ market power while retaining tractability.

The model also accounts for the possibility that the level of activity the platform can accommodate within each period may be limited, reflecting either time, resource, or facility constraints, as well as the possibility that the platform may provide auxiliary services, in limited supply, in addition to matching the agents.

The analysis first introduces and then studies the properties of a class of matching auctions that operate as follows. Upon joining the platform, agents select a membership status whose level determines the weight their bids receive in the subsequent auctions. At any period, agents are then asked to bid for each possible partner from the opposite side of the market. Each bilateral match then receives a “score” that depends on the involved agents’ reciprocal bids, on their membership status, and on the number of past interactions between the pair of agents. The matches with the highest nonnegative score are then implemented, up to capacity. As in the GSP auction, which can be seen as a simple static matching protocol, the payments the platform asks of each agent reflect the externalities the agent imposes on others due to the capacity constraint. Contrary to the GSP auction, however, such externalities are computed taking into account the value of experimentation and the dynamics of the agents’ informational rents. In particular, the platform may find it optimal to cross-subsidize certain interactions to generate information that can be used in future periods.

When match values evolve exogenously over time, the scores are *myopic*. The score assigned to each pair of agents coincides with the pair’s joint values, adjusted by (a) “handicaps” controlling for the informational rents the platform must leave to the agents, and (b) the platform’s own cost of implementing the matches, which reflect the relation-specific services provided by the platform, as discussed above. When, instead, match values evolve endogenously, as in the case of experimentation, or when agents have a preference for variety, the scores take the form of an *index*. As in other experimentation environments, such indexes summarize both the

current and future expected profitability of each match, taking into account the dynamics of (i) the agents' joint values, (ii) informational rents, and (iii) the platform's costs (or benefits) of implementing the matches.

In both cases, at all histories, including those off-path, in equilibrium, agents bid truthfully their myopic values for all partners. Bidding the myopic values is optimal for all agents because matching allocations under truthful bidding maximize continuation weighted surplus; that is, the sum of all agents' current and future payoffs, net of the platform's matching costs, and net of the agents' information rents. This property, together with the fact that the payments make each agent's continuation payoff proportional to the agent's contribution to continuation weighted surplus, then guarantees that each agent finds it optimal to stay in the mechanism and bid truthfully at all periods, irrespective of the agent's beliefs about other agents' current and past types, as well as of the implemented past matches. As a result, the proposed auctions are fully transparent: At the end of each period, all membership statuses and bids are disclosed.

One of the advantages of the proposed mechanisms is that they explicitly account for the value of experimentation, as advocated by many analysts and market designers, without, however, expecting the agents to adjust their bids to incorporate such values. The complexity of computing indexes which optimally control for the trade-offs between "experimentation" and "exploitation" is entirely on the platform's side. Once the agents understand the rules of the proposed mechanisms, it is in their interest to bid straightforwardly their myopic values in all periods. The proposed matching auctions allow for a high sensitivity of matches and prices to the evolution of the agents' preferences and information. As mentioned above, such sensitivity is particularly desirable in markets with a relatively small number of participating agents. Certain results, however, have implications also for larger markets, in which matching is more anonymous but where pricing is dynamic, responding to the evolution of aggregate market conditions. More generally, while, to the best of our knowledge, the specific auctions we propose are not used (yet) by any specific platform, we believe they capture the key trade-offs that platforms face when designing the dynamic algorithms disciplining the interactions among agents from various sides of the market, as reported in specialized outlets (see, e.g., the discussion in Chapter 6 in Moghaddam and Shimon (2016)).

The results also shed light on the value of regulating certain matching markets. The last few years have witnessed great interest (both from policy makers and academics) on how to regulate matching intermediaries (see, e.g., the articles "Online Platforms: Nostrums for Rostrums" and "Regulating Technology Companies: Taming the Beasts," *The Economist*, May 28, 2016). In markets in which all agents assign a nonnegative value to all interactions at every period, profit-maximizing intermediaries induce fewer and shorter interactions than what is efficient. Specifically, when the capacity constraint is not binding, a profit-maximizing intermediary matches each pair of agents for an inefficiently short period of time. When, instead, the capacity constraint is binding, certain interactions may last longer under profit maximization than under welfare maximization. However, the aggregate number of interactions in each period under profit maximization is always inefficiently low. Interestingly, the above conclusions need not extend to markets in which certain agents derive a negative payoff (equivalently, a payoff lower than their outside option) from interacting with certain other agents. In this case, profit maximization may result in an inefficiently large volume of matches, for any number of periods. The above conclusions have implications for how governments should subsidize certain platforms while taxing others, as well as for the costs and benefits of leaving matching markets unregulated.

The rest of the paper is organized as follows. We wrap up the introduction with a brief discussion of the most pertinent literature. Section 2 presents an example of a simple dynamic matching environment whose role is to introduce the reader to the main ideas developed later in the paper. Section 3 describes the full model. Section 4 introduces the matching auctions. Section 5 derives equilibrium properties of the proposed mechanisms. Section 6 identifies a subclass of matching auctions that are profit-maximizing. Section 7, instead, identifies a class of auctions that are welfare-maximizing and contrasts matching dynamics under profit maximization with their counterparts under welfare maximization. Section 8 concludes. Proofs omitted in the main text are either in the Appendix at the end of the document, or in the article’s supplementary material.

1.1 Related Literature

Dynamic matching. Most of the recent literature on centralized dynamic matching focuses on markets without transfers, in which agents are matched once, with dynamics stemming from the arrival and departure of agents to and from the market. In the context of kidney exchange, Ünver (2010) studies optimal mechanisms for two-way and multi-way exchanges, minimizing total waiting costs, in a market with stochastic arrivals of donors and recipients. Optimal dynamic matching is also the focus of Anderson, Ashlagi, Gamarnik, and Kanoria (2015), Baccara, Lee, and Yariv (2015), Akbarpour, Li, and Oveis Gharan (2016), and Herbst and Schickner (2016). A key trade-off in such environments is between avoiding waiting costs and waiting for the market to thicken. A related strand of papers study the assignment of objects through waiting lists (see Leshno (2015), Thakral (2015), Bloch and Cantala (2016), and Schummer (2016) for recent developments).⁹

As anticipated above, the key difference with respect to this literature is that, in the present paper, agents change partners in response to changes in actual or perceived valuations, as opposed to the arrival and departure of agents to and from the market.

Matching with transfers. The paper is also related to the literature on profit-maximization in matching markets with private information and transfers. Damiano and Li (2007) and Johnson (2013) consider a *one-to-one* matching environment where the intermediary faces asymmetric information about agents’ vertical characteristics responsible for match values. Board (2009) studies the problem of a profit-maximizing platform (e.g., a school) that can induce agents to self-select into mutually exclusive groups (e.g., classes). These papers derive conditions on primitives for a profit-maximizing intermediary to induce positive assortative matching. Gomes and Pavan (2016) study many-to-many matching in a setting where agents differ both in their consumer value (willingness-to-pay) and input value (salience) and where matching is non-partitional.¹⁰

Matching in all of these papers is static. In contrast, the present paper considers dynamic matching in an environment in which match values evolve, either exogenously, or endogenously, over time. Fershtman and Pavan (2017) offers an example of a dynamic auction for a matching environment in which each agent perfectly learns the attractiveness of each partner after the first interaction (we labelled such auctions “Pandora auctions” in reference to the Pandora’s box problem studied in the operations research literature). That paper builds on the present one. In particular, all the results in that paper are special cases of the general results developed

⁹See also Damiano and Lam (2005), Kurino (2009), and Doval (2015) for appropriate stability notions in dynamic matching environments.

¹⁰See also Gomes and Pavan (2017) for a static model of mediated matching in which agents’ preferences combine elements of vertical differentiation with elements of horizontal differentiation.

in the present paper. Section 2 below offers another example of an environment in which the general results developed in the present paper can be applied to. Contrary to Fershtman and Pavan (2017), in the example in Section 2, agents keep refining their beliefs about the attractiveness of their partners over time till they find out that certain interactions are unprofitable. The example also illustrates the role of a certain separability condition in settings in which the platform can accommodate more than a single match, but not all matches, in each period.

Position and scoring auctions. In procurement auctions, scoring rules are often used to aggregate the various dimensions of the sellers' offers (price, product design, delivery time, etc.). See, for example, Che (1993), Asker and Cantillon (2008). Our matching auctions share with this literature the idea that the desired allocations can be induced through an appropriate design of the scoring rules governing the auctions. However, while the above literature focuses on static settings, our scoring rules are for dynamic environments in which preferences evolve over time. Another difference is that our scores aggregate the preferences of different agents from different sides of the market, whereas in the procurement auctions in the aforementioned literature the scores aggregate the various dimensions of each seller's own offer.

Another related literature studies auctions for sponsored links by search engines and for online ads by ad exchanges. For example, Varian (2007), Edelman, Ostrovsky and Schwarz (2007), and Gomes and Sweeney (2014) study the properties of the GSP auction, used by online search engines to allocate ads,¹¹ while Mansour, Muthukrishnan, and Nisan (2012) study auctions used by ad exchanges to match advertisers with content providers. Our matching auctions are relevant also for these markets, modulo the fact that, in online search, searchers typically do not pay for matches (or, more precisely, the currency used for the services they receive is the release of their privacy).¹² Notwithstanding this qualification, one of the advantages of our auctions is that they explicitly account for the value of generating information that can be used in future periods.

Platforms. Markets where agents purchase access to other agents are the focus of the literature on two-sided markets (see Rysman (2009) for a survey, and Weyl (2010), Bedre-Defolie and Calvano (2013), Lee (2013), and Jullien and Pavan (2017) for some of the recent developments). This literature restricts attention to a single network or to mutually exclusive networks. In contrast, the present paper allows for general matching rules and for more flexible payoff structures. In particular, it does not restrict agents' willingness to pay to coincide with their attractiveness. Most importantly, it focuses on a dynamic environment in which match values change over time. Cabral (2011) also considers a dynamic model with network effects but in which values are constant over time.

Screening under experimentation. The paper is also related to the literature on screening agents who experiment about the value of a good. In particular, the mechanisms proposed here can be seen as the matching analogs of the bandit auctions of Pavan et al. (2014) and Kakade et al. (2013) for the sale of an indivisible physical object. See also Bergemann and Valimaki (2006) for a survey on bandit problems in economics.

Mechanism design. From a methodological standpoint, we draw from recent developments in the dynamic mechanism design literature. In particular, the conditions for incentive compatibility in the present paper adapt

¹¹See also Athey and Ellison (2011), Börgers, Cox, Pesendorfer and Petricek (2013), and Gomes (2014).

¹²In ad exchange auctions, instead, it is becoming customary for content providers to specify dynamic reservation prices for different types of ads, which is a form of dynamic bidding. In these auctions, both sides thus dynamically bid for the matches.

to the environment under examination results in Bergemann and Valimaki (2010) and Pavan, Segal, and Toikka (2014). See also Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Board (2007), and Eso and Szentes (2007), for some of the earlier contributions, Borgers (2015), Bergemann and Pavan (2015), Bergemann and Valimaki (2017), and Pavan (2017) for general overviews of this literature, and Bergemann and Strack (2015) for recent developments in continuous time. In most of these models, the agents’ private information is exogenous. In contrast, the agents’ private information is endogenous in the present paper, in the bandit single-unit auction of Pavan et al. (2014), in the dynamic virtual pivot mechanism of Kakade et al. (2013), and in the taxation model of Makris and Pavan (2017).

Particularly related is the strand of the dynamic mechanism design literature that investigates how to implement dynamically efficient allocations in settings in which the agents’ types evolve over time, thus extending the Vickrey–Clarke–Groves (VCG) and d’Aspremont–Gérard-Varet (AGV) results from static to dynamic settings (see, for example, Bergemann and Valimaki, 2010, Athey and Segal, 2013, and the references therein). The payments in the auctions we propose here are determined by a pricing formula similar to those in Bergemann and Valimaki (2010) and in Kakade et al. (2013), adapted to the fact that the private information the agents receive in each period is multi-dimensional, and the fact that the allocation rule under profit maximization need not be efficient.

Another stream of the dynamic mechanism design literature considers both efficient and profit-maximizing mechanisms in settings where the agents’ private information is static, but where the dynamics originate in the arrival and/or departure of objects and agents to and from the market (for an overview of this literature, see Bergemann and Said (2011), and Gershkov and Moldovanu (2014)).

2 Example

The intermediary is a media platform owning two outlets. Each outlet can be used to match content providers with advertisers. To illustrate the key ideas in the simplest possible way, assume there are two content providers and two advertisers. Let the content providers belong to side A and the advertisers to side B and denote by $N_A = \{1, 2\}$ and $N_B = \{1, 2\}$ the respective sets of agents. Each outlet can be used to match at most one content provider with one advertiser. However, each advertiser and each content provider can be present on multiple outlets (i.e., matching is many-to-many). In order to preserve heterogeneity across the two outlets, the platform must guarantee that the same advertiser-content provider pair is not present on both outlets at the same time.

The “agents” (content providers and advertisers), as well as the platform, are uncertain about the profitability of each potential match. Importantly, content providers and advertisers learn from previous interactions privately. For example, a content provider may learn that having its content displayed together with a certain type of advertisement reflects poorly on its sponsors. Similarly, an advertiser may learn that having its ads displayed with a certain type of content does not attract the type of audience it wishes. Because of such uncertainty, the platform values retaining flexibility and being able to change the composition of the outlets in response to variations in the perceived profitability of the matches. Because the latter is the agents’ private information, the platform uses auctions to determine the composition of the matches over time.

From the perspective of each content provider $i = 1, 2$, the match with each advertiser $j = 1, 2$ may either be good, $\omega_{ij}^A = g$, in which case it delivers a flow profit equal to $u_{ij}^A > 0$, or bad, $\omega_{ij}^A = b$, in which case it delivers a flow loss equal to $-u_{ij}^A$.¹³ A similar structure describes the profitability of the matches from the perspective of the advertisers.¹⁴ Hence, a match that is good from the perspective of one agent may be bad from the perspective of the other.

The match values u_{ij}^k are commonly known by all players, whereas the qualities of the matches ω_{ij}^k are unknown to both the agents and the platform. What is commonly known is that ω_{ij}^k are independently drawn. Let $\eta_{ij0}^k = \Pr(\omega_{ij}^k = g)$ denote the initial probability that all players assign to the quality of the match between i and j being good for the agent on side $k \in \{A, B\}$. For simplicity, we assume here that such a probability is commonly known; the general model, however, accommodates for the possibility that it is i 's private information. Likewise, let η_{ijt}^k denote the probability that $\omega_{ij}^k = g$ at the beginning of period t and

$$v_{ijt}^k = \eta_{ijt}^k u_{ij}^k - (1 - \eta_{ijt}^k) u_{ij}^k$$

the expected flow payoff of the match between i and j , for the agent on side $k \in \{A, B\}$. Importantly, both η_{ijt}^k and v_{ijt}^k depend on the outcomes of past interactions, as explained below.

After each interaction with agent j , agent i from side k receives a signal $\xi_{ij}^k \in \{g, b\}$ about the quality of the match with j drawn independently across matches according to the following structure. When $\omega_{ij}^k = g$, $\xi_{ij}^k = g$ almost surely. When, instead, $\omega_{ij}^k = b$,

$$\xi_{ij}^k = \begin{cases} g & \text{w.p. } \mu_{ij}^k \\ b & \text{w.p. } 1 - \mu_{ij}^k. \end{cases}$$

That is, a bad signal perfectly reveals that the match is unprofitable, whereas a good signal may be a "false positive" in the sense that it is also consistent with the quality of the match being bad. In the terminology of the experimentation literature, the above structure thus corresponds to the "no-news-is-good-news" case.¹⁵ Importantly, all the signals are drawn independently across pairs, conditional on the state $\omega = (\omega_{ij}^k)_{i,j=1,2}^{k=A,B}$.

The platform's flow profits consist of the payments received from the agents, net of the match-specific per-period costs $c_{ij} \in \mathbb{R}$. The latter can be positive or negative depending on the match. For example, while certain matches may help promote the platforms' outlets, others may be detrimental to its position in the media market. Each agent's flow payoff is equal to the sum of the expected values of the matches the agent is involved in, net of the payments to the platform.

Time is discrete and infinite, and denoted by $t = 1, 2, \dots$, and all players (the platform, as well as all the

¹³The superscripts $k = A, B$ denote the side of the agent whose perspective is taken. Also note that it is only for simplicity that we assume that the loss of a bad match equals the profit from a good one. The model in the next section dispenses with most of the simplifying assumptions in this example.

¹⁴Denoting by $\omega_{ij}^B \in \{g, b\}$ the quality of the match with content provider i , the flow payoff that advertiser j obtains from being matched with content provider i is thus equal to u_{ij}^B if $\omega_{ij}^B = g$ and is equal to $-u_{ij}^B$ if $\omega_{ij}^B = b$.

¹⁵Posterior beliefs about the quality of the matches then evolve according to the following process: if (i, j) were not matched in period $t - 1$, then $\eta_{ijt}^k = \eta_{ijt-1}^k$. If, instead, the pair was matched and i received at least one b signal in one of the past interactions, then $\eta_{ijt}^k = 0$. Finally, if the pair was matched at $t - 1$ and all the signals received in previous interactions, including the period- $(t - 1)$ one, were g , then

$$\eta_{ijt}^k = \frac{\eta_{ijt-1}^k}{\eta_{ijt-1}^k + (1 - \eta_{ijt-1}^k)\mu_{ij}^k} > \eta_{ijt-1}^k.$$

agents) maximize the expected discounted sum of their flow payoffs using the common discount factor $\delta \in (0, 1)$. All players' outside options yield a flow payoff equal to zero.

The environment described above is a stylized version of the more general environment developed in the next section. In particular, while in this example agents possess no private information prior to entering into negotiations with the platform and perfectly learn the quality of each match after receiving a bad signal, in the general model, some agents may possess private information prior to contracting with the platform, and learning may continue for arbitrary long horizons.¹⁶

As anticipated above, because the outcomes of previous interactions are the agents' private information, the platform uses a dynamic matching protocol to determine which matches to implement over time. In each period $t = 1, 2, \dots$, each agent is asked to submit a bid b_{ij}^k for each possible match. The bids can be either positive or negative reflecting the fact that while certain matches may be perceived as profitable, others may be seen as unprofitable. The bids submitted by the agents in period t , along with those submitted in previous periods, then determine the matches implemented in period t as well as the period- t payments. Because in this example the agents possess no private information at the time they first join the auction, the platform can extract the entire surplus from the agents. The matches that the platform induces over time are thus those that maximize the expected discounted sum of all agents' gross payoffs, net of the platform's costs. The analysis in the subsequent sections identifies conditions under which the matches that maximize the platform's objective (e.g., profits, or welfare) can be induced by assigning to each match $(i, j) \in N_A \times N_B$ a "score" that depends only on information that pertains to that match (i.e., current and past bids submitted by agents i and j for the match (i, j)) and then selecting in each period the matches with the highest nonnegative score. If all scores are negative, no match is formed. While a "scoring rule" may look natural, note that many dynamic allocation rules cannot be described by comparing the scores for the individual matches. In addition to the matching rule, the dynamic matching protocol must also specify a payment rule, mapping the bids into the payments asked to the agents over time.

In the context of the simple environment considered in this example, given the value of experimenting with different matches, a natural candidate for the profit-maximizing mechanism is a dynamic scoring rule where the scores are linked to the Gittins indexes from the experimentation literature (see, for example, Gittins 1979). Such indexes combine the myopic expected profitability of the matches based on past and current information, with the value of generating new information that can be used in subsequent periods. A matching auction that induces the agents to bid according to the indexes they assign to the various matches, however, has two problems. First, it may be difficult, in practice, for the agents to compute such indexes. Second, and more importantly, the private value that agent i from side k may assign to experimenting with agent j from side $l \neq k$ need not coincide with either the value that agent j may assign to experimenting with i , or with the platform's value. For example, while i may have already learned that the match with j is unprofitable, j may face residual uncertainty about the quality of his match with j . A simple solution would then consist in designing an auction that induces all agents to simply bid myopically their expected flow payoffs and then use scores for the individual matches that take into account the value of cross-subsidizing certain interactions. The score assigned to each match (i, j) would then take the form of the Gittins index for the arm whose flow rewards

¹⁶With the notation in the general model, the example corresponds to a situation where $\theta_i^k = 1$, $i = 1, 2$, $k = A, B$.

correspond to the sum of the involved agents' flow payoffs $v_{ijt}^A + v_{ijt}^B$, net of the platform's costs c_{ij} . Provided the matching allocations induced under such rule maximize the sum of all players' expected continuation payoffs after all histories, the agents could then be incentivized to bid truthfully by linking the flow payments to the agents' contributions to expected continuation surplus, using, for example, the dynamic pivot payment rule of Bergemann and Valimaki (2010).

The above mechanism, however, faces two difficulties. First, in general, the matches that maximize the platform's ex-ante payoff need not coincide with those that maximize the sum of all players' continuation payoffs at all histories. This is the case, for example, when certain agents possess private information prior to joining the mechanism and the platform's objective is the maximization of its profits. Second, even when the platform maximizes total surplus, the dynamics induced by selecting the matches that have the highest non-negative scores need not maximize the expected continuation total surplus at all histories. The reason is that, in each period, more than a single match may be formed. The problem faced by the platform is thus akin to a multiarmed bandit problem where, in each period, more than a single arm is activated. As is well known from the experimentation literature, in general, such problems do not admit a simple solution taking the form of an index rule.¹⁷

An alternative possibility would be to assign an index to each possible pair of matches and then select, in each period, the two pairs for which the index is the highest. The problem with such a construction is that, because certain agents are involved in more than a single match, the flow rewards of such aggregate arms are correlated. With correlated rewards, in general, an index policy fails to be optimal.

A special case of interest where the profit-maximizing mechanism takes the form of a dynamic matching auction, selecting in each period the matches with the highest bilateral scores with the score for each match (i, j) taking the form of a Gittins index for the arm with flow rewards $v_{ijt}^A + v_{ijt}^B - c_{ij}$, is when the environment satisfies an appropriate "separability" condition. Let M denote the number of matches the platform can implement per period. In this example, $M = 2$. The separability condition, formally defined in Section 5, requires that, in each period, any two matches for which the historical infimum of the Gittins index is the M -th and $(M + 1)$ -th highest have sufficiently different historical minimal values of the Gittins indexes, unless one of the two historical values is negative, in which case no restriction is imposed. Note that the property imposes no restriction on the upside risk of each index, only the downside.

To illustrate, let $\lambda_{ij} | b_{ijt}^A, b_{ijt}^B, m$ denote the process over the sum of agents i and j 's flow payoffs when the pair interacted m times prior to period t and the current expected flow payoffs are equal to $(v_{ijt}^A, v_{ijt}^B) = (b_{ijt}^A, b_{ijt}^B)$, where b_{ijt}^A and b_{ijt}^B are the two agents' period- t bids for their match. The period- t index is then defined by

$$S_{ijt} \equiv \max_{\tau} \left\{ \frac{\mathbb{E}^{\lambda_{ij} | b_{ijt}^A, b_{ijt}^B, m} \left[\sum_{s=t}^{\tau} \delta^{s-t} \left(\tilde{v}_{ijs}^A + \tilde{v}_{ijs}^B - c_{ij} \right) \right]}{\mathbb{E}^{\lambda_{ij} | b_{ijt}^A, b_{ijt}^B, m} \left[\sum_{s=t}^{\tau} \delta^{s-t} \right]} \right\},$$

where τ is a stopping time. For example, when, for each $(i, j) \in N_A \times N_B$, $u_{ij}^A = u_{ij}^B = u_{ij}$, $\eta_{ij0}^A = \eta_{ij0}^B = \eta_{ij0} \geq \frac{1}{2}$,

¹⁷Example 1 in the Supplementary Material provides an illustration of why index rules may be suboptimal. See also Bergemann and Valimaki (2001) for an additional example.

and $\mu_{ij}^A = \mu_{ij}^B = \mu_{ij}$, the period- t Gittins index for each match (i, j) is equal to (the proof is in the Appendix)

$$S_{ijt} = 2u_{ij} \left(\frac{\frac{\eta_{ij0}^2}{1-\delta} - \frac{\mu_{ij}^{2m}(1-\eta_{ij0})^2}{1-\delta\mu_{ij}^2}}{\frac{\eta_{ij0}^2}{1-\delta} + \frac{2\eta_{ij0}(1-\eta_{ij0})\mu_{ij}^m}{1-\delta\mu_{ij}} + \frac{\mu_{ij}^{2m}(1-\eta_{ij0})^2}{1-\delta\mu_{ij}^2}} \right) - c_{ij},$$

when both b_{ijt}^A and b_{ijt}^B are positive. When, instead, one of the two bids is negative and the other is positive, the index is equal to

$$S_{ijt} = -2u_{ij} \left(\frac{\frac{\mu_{ij}^m(1-\eta_{ij0})}{1-\delta\mu_{ij}}}{\frac{\eta_{ij0}}{1-\delta} + \frac{\mu_{ij}^m(1-\eta_{ij0})}{1-\delta\mu_{ij}}} \right) - c_{ij}.$$

Finally, when both bids are negative, then $S_{ijt} = -1$. Note that, because c_{ij} can be negative, a single negative bid can still be compatible with the index being positive.

Now let

$$\underline{S}_{ijt} \equiv \inf_{s \leq t} \{S_{ijs}\}$$

denote the historical infimum of the index corresponding to the match (i, j) . Separability then amounts to the requirement that, in each period t , any two matches (i, j) and (i', j') for which \underline{S}_{ijt} and $\underline{S}_{i'j't}$ are, respectively, the second and third highest, are such that, if $\underline{S}_{ijt} > \underline{S}_{i'j't} > 0$, then $\underline{S}_{ijt}(1-\delta) \geq \underline{S}_{i'j't}$. As an illustration, observe that, in the example, the condition is satisfied when $(u_{11}, \mu_{11}, \eta_{110}, c_{11}) = (\frac{1}{5}, \frac{4}{5}, \frac{4}{5}, 0)$, $(u_{12}, \mu_{12}, \eta_{120}, c_{12}) = (2, \frac{1}{5}, \frac{1}{2}, -\frac{1}{100})$, $(u_{21}, \mu_{21}, \eta_{210}, c_{21}) = (2, \frac{1}{5}, \frac{1}{2}, 0)$, $u_{22} = c_{22} = 0$, and $\delta = \frac{4}{5}$. In addition to satisfying the aforementioned separability condition, the numerical values have been selected to generate dynamics with the following interesting properties: (a) during the initial phase, the platform favors matches with a lower expected flow joint match value, (b) after the arrival of negative news, a match may be interrupted momentarily but then reactivated after another match also receives negative news, (c) whether a match that is suspended momentarily is eventually reactivated may depend on the number of periods the agents involved in the match interacted prior to receiving negative news.

In particular, note that the example has been constructed so that the match between the second content provider and the second advertiser is unprofitable. Furthermore, the matches (1,2) and (2,1) share the same primitives, except for the fact that the platform derives a direct benefit equal to $1/100$ from the former match and no benefit from the latter match. Finally, the match (1,1) is more likely to be good than the matches (1,2) and (2,1), although, conditional on being good, the former match yields a lower payoff to each of the two agents and less informative signals than the latter matches.

The following dynamics are then sustained under both the welfare- and the profit-maximizing auctions. Irrespective of the submitted bids, in period one, the platform implements the matches (1,2) and (2,1), despite these two matches being myopically inferior to (1,1). These two matches are continued until one of the agents involved in such matches submits a negative bid, thus revealing he received a negative signal.¹⁸ When this happens, the match involving such an agent is replaced by the match (1,1). If, within the same period, both the match (1,2) and (2,1) contain negative bids, then (2,1) is replaced with (1,1), whereas (1,2) is retained. Once

¹⁸In this simple example, the only relevant information contained in the bids is their sign, which indicates whether or not an agent received a negative signal. In the general model, the magnitude of the bids also plays a role.

one of the agents involved in the matches (1,1) or (2,1) submits a negative bid, that match is never implemented again. Finally, the match (1,2) may be interrupted upon one of the agents involved in the match submitting a negative bid but later resumed (with the agent submitting the negative bid compensated for his disutility) until the other agent involved in the match also submits a negative bid, at which point the match is permanently dismissed. Importantly, whether or not the match (1,2) is reactivated depends on the number of periods during which both agents involved in the match submitted positive bids before one of the two submitted a negative bid. In particular, the match is reactivated after being temporarily suspended provided the number of periods in which the two agents both submitted positive bids is greater than two.

The above example is stylized in many ways: there are two agents on each side, certain matches are known to be unprofitable, agents perfectly learn about the profitability of a match after receiving negative news, and, importantly, agents possess no private information prior to joining the platform. The last assumption implies that the dynamics under profit-maximization coincide with those under welfare-maximization. In the general model, the above restrictions are dispensed with. Among other things, the analysis identifies conditions under which profit maximization leads to excessively short experimentation, as well as conditions under which experimentation under profit maximization lasts longer than under welfare maximization.

3 The Environment

Agents, matches, and preferences

A matchmaking platform mediates the interactions among agents from two sides of a market, A and B . There are $n_A \in \mathbb{N}$ agents on side A and $n_B \in \mathbb{N}$ agents on side B , with $N_A \equiv \{1, \dots, n_A\}$ and $N_B \equiv \{1, \dots, n_B\}$ denoting the corresponding sets of agents on the two sides. Time is discrete, indexed by $t = 0, 1, \dots, \infty$. Agents live for infinitely many periods and can change partners infinitely many times.

Below, we describe various features of the environment from the perspective of a generic agent from side A . A similar description applies to side B .

The flow period- t payoff that agent $i \in N_A$ from side A derives from being matched to agent $j \in N_B$ from side B is given by

$$v_{ijt}^A(\theta_i^A, \varepsilon_{ijt}^A) = \theta_i^A \cdot \varepsilon_{ijt}^A. \quad (1)$$

The parameter θ_i^A is time- and match-invariant and captures the overall importance that agent i assigns to interacting with agents from the opposite side of the market. That is, θ_i^A parametrizes the value that agent i assigns to a generic interaction with an agent from side B , prior to conditioning on the specific profile of the latter agent. The parameter ε_{ijt}^A , instead, is match-specific and time-variant and captures the attractiveness of agent j from side B in the eyes of agent i . These match-specific values evolve over time, reflecting the change in the agents' true, or perceived, attractiveness. They can either represent the evolution of the agents' beliefs about fixed, but unknown, match qualities, or variations in attractiveness triggered by stochastic changes in the environment. Hereafter we refer to θ_i^A as the agent's "*vertical type*" and to $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N_B}$ as the profile of the agent's period- t "*horizontal types*". We refer to $v_{it}^A \equiv (v_{ijt}^A)_{j \in N_B}$ as the agent's period- t "*match values*".

Both the agents' vertical and horizontal types are their own private information. Each agent's vertical

type θ_i^A is drawn from an absolutely continuous cumulative distribution function F_i^A with density f_i^A strictly positive over $\Theta_i^A = [\underline{\theta}_i^A, \bar{\theta}_i^A]$, with $\underline{\theta}_i^A > 0$. Vertical types are drawn independently across agents and from the horizontal types $\varepsilon \equiv (\varepsilon_{ijt}^k)_{t=1, \dots, \infty}^{(i,j) \in N_A \times N_B, k=A, B} \in \mathcal{E}$. Importantly, while we restrict the agents' vertical types to be nonnegative, we allow the horizontal types to be negative, reflecting the possibility that an agent may derive a negative utility from interacting with certain agents from the opposite side.¹⁹

Agents learn their vertical types prior to entering into negotiations with the platform and then learn the horizontal types gradually over time, with each ε_{ijt}^A learned at the beginning of period t , $t = 1, \dots, \infty$.

For any $t \geq 1$, and any pair of agents $(i, j) \in N_A \times N_B$, let $X_{ijt} \equiv \{0, 1\}$, with $x_{ijt} = 1$ in case the pair is matched in period t , and with $x_{ijt} = 0$ in case the pair is unmatched. Matching is many-to-many, meaning that the same agent may be matched to multiple agents from the opposite side.

All agents are expected-utility maximizers and maximize the expected discounted sum of their flow payoffs using the common discount factor $\delta \in (0, 1]$. Let $p_t \equiv (p_{it}^k)_{i \in N_k, k=A, B}$ denote the payments collected by the platform from the two sides of the market in period t , and $p \equiv (p_t)_{t=0}^\infty$ an entire sequence of payments. Note that, while the matching starts in period one, we allow the platform to start collecting payments from the agents in period zero, after the agents have observed their vertical types, but before they observe their horizontal partner-specific types. This assumption is motivated by the idea that, in many markets of interest, at the time the agents join the platform, they do not know yet the specific profile of the agents who join from the opposite side. In other words, agents learn about agents from the opposite side only after "getting on board". Also note that payments are allowed to be negative, reflecting the possibility that (a) certain agents may dislike certain interactions and ask to be compensated, or (b) even if all agents like interacting with all other agents, the platform may want to cross-subsidize certain interactions.

The (Bernoulli) payoff function for each agent i from side A is given by

$$U_i^A = \sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} x_{ijt} v_{ijt}^A - \sum_{t=0}^{\infty} \delta^t p_{it}^A. \quad (2)$$

The platform's (Bernoulli) payoff function is the discounted sum of the payments collected from the two sides of the market, net of possible costs of implementing the matches:

$$U_0 = \sum_{t=0}^{\infty} \delta^t \left(\sum_{i \in N_A} p_{it}^A + \sum_{j \in N_B} p_{jt}^B \right) - \sum_{t=1}^{\infty} \delta^t \left(\sum_{i \in N_A} \sum_{j \in N_B} c_{ijt}(x_{ij}^{t-1}) \cdot x_{ijt} \right),$$

where $c_{ijt} \in \mathbb{R}$ is the period- t "cost" of matching the pair (i, j) , with $x_{ij}^{t-1} \equiv (x_{ijs})_{s=1}^{t-1} \in X_{ij}^{t-1} \equiv \prod_{s=1}^{t-1} X_{ijs}$ denoting the history of past interactions between the pair (i, j) . We allow the costs to take on negative values so as to capture the possibility that, in certain applications, the platform may derive a direct benefit from certain matches, for example because they can be used to advertise the platform's matching capabilities. As anticipated in the Introduction, these costs also incorporate all the auxiliary services that the platform may provide to the agents, over and above connecting one agent with the others.

¹⁹Under the assumed multiplicative structure $v_{ijt}^A = \theta_i^A \cdot \varepsilon_{ijt}^A$, allowing the vertical types to also take on negative values would introduce confusion given that the horizontal types ε_{ijt}^A are already allowed to take on negative values.

Evolution of match values

The match values v_{ijt}^A that agent i from side A derives from interacting with agent j from side B are correlated over time, both through the fully persistent vertical component θ_i^A and through the partially persistent horizontal components ε_{ijt}^A . We assume that the latter evolve over time according to the following process. For each pair $(i, j) \in N_A \times N_B$, and each period $t \geq 1$, ε_{ijt}^A may take any value in $\mathcal{E}_{ijt}^A \subseteq \mathbb{R}$. While not essential to the results, we find it convenient to think of \mathcal{E}_{ijt}^A either as the entire real line, or as a compact and connected subset of it.

The evolution of agent i 's horizontal types is governed by a collection of match-specific kernels $G_{ij}^A \equiv (G_{ijt}^A : \varepsilon_{ijt}^A \times \mathcal{E}_{ijt-1}^A \times X_{ij}^{t-1} \rightarrow [0, 1])_{t=1}^\infty$, with $X_{ij}^0 \equiv \emptyset$. We then let $G \equiv (G_{ij}^k)_{(i,j) \in N_A \times N_B}^{k=A,B}$ denote the complete collection of kernels for all matches. The interpretation of these kernels is the following. Each period-1 horizontal type ε_{ij1}^A is drawn from the cdf G_{ij1}^A . In each subsequent period $t > 1$, the period- t type ε_{ijt}^A is drawn from the cdf $G_{ijt}^A(\varepsilon_{ijt}^A | \varepsilon_{ijt-1}^A, x_{ij}^{t-1})$, where ε_{ijt-1}^A is the agent's horizontal match-specific type from the preceding period, and where x_{ij}^{t-1} is the history of past interactions between the pair (i, j) . Importantly, while the support of each kernel G_{ijt}^A is a subset of \mathcal{E}_{ijt}^A , we allow for the possibility that, for certain histories $(\varepsilon_{ijt-1}^A, x_{ij}^{t-1})$, it is a strict subset of \mathcal{E}_{ijt}^A .

To guarantee that the expected payoff of each agent is well defined at all histories, and satisfies a certain envelope formula (more below), we assume that, for all $i \in N_A$, there exists a constant $E_i^A > 0$ such that, for any sequence of matches x , $\mathbb{E} \left[\sum_{t=1}^\infty \delta^t \sum_{j \in N_B} |\varepsilon_{ijt}^A| \cdot x_{ijt} \right] \leq E_i^A$, where the expectation is taken with respect to the distribution over \mathcal{E} generated by the kernels G , under the matches x .

We will focus on two environments that capture different features of various dynamic matching markets with private information. The first environment is one in which the evolution of the agents' match values and of the platform's costs is exogenous. The second environment is one in which the match values and the platform's costs depend on past interactions, with properties reflecting either private experimentation or preferences for variety.

- **Exogenous processes.** For all $(i, j) \in N_A \times N_B$, $t > 1$, and all $(\varepsilon_{ijt-1}^A, \varepsilon_{ijt}^A) \in \mathcal{E}_{ijt-1}^A \times \mathcal{E}_{ijt}^A$, $G_{ijt}^A(\varepsilon_{ijt}^A | \varepsilon_{ijt-1}^A, x_{ij}^{t-1})$ is invariant in x_{ij}^{t-1} . Furthermore, c_{ijt} does not depend on x_{ij}^{t-1} .

- **Endogenous processes.** For any $(i, j) \in N_A \times N_B$, the following properties hold: (i) whenever $x_{ijt-1} = 1$, the dependence of $G_{ijt}^A(\varepsilon_{ijt}^A | \varepsilon_{ijt-1}^A, x_{ij}^{t-1})$ on x_{ij}^{t-1} is only through $\sum_{s=1}^{t-1} x_{ijs}$; (ii) whenever $x_{ijt-1} = 0$, G_{ijt}^A is a Dirac measure at $\varepsilon_{ijt}^A = \varepsilon_{ijt-1}^A$, i.e., $G_{ijt}^A(\varepsilon_{ijt}^A | \varepsilon_{ijt-1}^A, x_{ij}^{t-1}) = \mathbf{1}_{\{\varepsilon_{ijt}^A \geq \varepsilon_{ijt-1}^A\}}$; (iii) there exists a sequence $(\omega_{ijs}^A)_{s=1}^\infty \in \mathbb{R}^\infty$ drawn from an exogenous distribution, such that, for any number R_{ij} of past interactions between agent $i \in N_A$ and agent $j \in N_B$, the period- t match value ε_{ijt}^A is given by a deterministic function of $(\omega_{ijs}^A)_{s=1}^{R_{ij}}$, uniformly over t ; (iv) the cost $c_{ijt}(x_{ij}^{t-1})$ depends on x_{ij}^{t-1} only through $\sum_{s=1}^{t-1} x_{ijs}$.

In the case of endogenous processes, the above assumptions imply the following properties: (1) agents' match-specific values change only upon interacting with partners; (2) The processes governing the agents' match values are Markov time-homogeneous — if agents $(i, j) \in N_A \times N_B$ are matched in period $t - 1$, the distribution of ε_{ijt}^A depends only on ε_{ijt-1}^A and the number of past interactions between the pair (i, j) . That the platform's cost of matching a pair of agents depends on the number of times the pair interacted in the past reflects the idea that, in most markets of interest, costs decrease with the number of past interactions. A

case of particular interest is when the cost vanishes after a link between the pair is created, i.e., after the first interaction.

Example 1 (Gaussian learning) Suppose that every agent $i \in N_A$ from side A derives a constant utility u_{ij}^A from interacting with each agent $j \in N_B$ from side B , and that this utility is unknown to the platform and to all agents. Agent i starts with a prior belief that $u_{ij}^A \sim N(\varepsilon_{ij1}^A, \tau_{ij}^A)$, where the variance τ_{ij}^A is common knowledge but where the initial prior mean ε_{ij1}^A is the agent's private information. The agent's prior mean ε_{ij1}^A is drawn from a distribution G_{ij1}^A . Each time agent i is matched to agent j , agent i receives a conditionally i.i.d. private signal $\xi_{ij}^A \sim N(u_{ij}^A, \vartheta_{ij}^A)$ about u_{ij}^A and updates his expectation of u_{ij}^A using standard projection formulae.²⁰ Let ε_{ijt}^A be agent i 's posterior mean about u_{ij}^A in period t . Such environment satisfies the assumptions of the above endogenous-processes model. In particular, when $x_{ijt-1} = 1$, agent i 's posterior mean ε_{ijt}^A about u_{ij}^A is drawn from a (Gaussian) distribution G_{ijt}^A that depends only on the agent's mean ε_{ijt-1}^A in the previous period and the number of past interactions $\sum_{s=1}^{t-1} x_{ijs}$ with agent j . \square

Example 2 (preference for variety) Suppose that the value each agent $i \in N_A$ derives from interacting with each agent $j \in N_B$ decreases (possibly stochastically) with the number of past interactions with agent j . Precisely, for all $t \geq 1$, all $j \in N_B$, all $(\varepsilon_{ijt-1}^A, \varepsilon_{ijt}^A)$, $G_{ijt}^A(\varepsilon_{ijt}^A \mid \varepsilon_{ijt-1}^A, x_{ijs}^{t-1})$ is non-decreasing in $\sum_{s=1}^{t-1} x_{ijs}$. This assumption captures the idea that agents gradually lose interest in partners with whom they interacted already. Alternatively, agents may have a fixed demand for interactions with each partner, or each agent may provide only up to a fixed number of services to each of his partners. Such environments are special cases of the endogenous-processes model. For example, the case in which each agent has a fixed demand for interactions with each partner can be captured by assuming that, for each agent $i \in N_A$, each partner $j \in N_B$, there exists a number $\alpha_{ij}^A \in \mathbb{N}$ such that, whenever $\sum_{s=1}^{t-1} x_{ijs} > \alpha_{ij}^A$, $G_{ijt}^A(\varepsilon_{ijt}^A \mid \varepsilon_{ijt-1}^A, x^{t-1}) = 1$ for all $\varepsilon_{ijt}^A \geq 0$. \square

Remarks. Given the matches x , agent i 's horizontal types $\varepsilon_{ij}^A \equiv (\varepsilon_{ijt}^A)_{t=1, \dots, \infty}$ for interacting with agent j from side B are drawn independently from the values $\varepsilon_{ij'}^A = (\varepsilon_{ij't}^A)_{t=1, \dots, \infty}$ the agent derives from interacting with any other agent $j' \in N_B$. When processes are exogenous, this assumption does not play any role and can be dispensed with. With endogenous processes, instead, such an assumption facilitates the characterization of the matching dynamics under both the profit-maximizing and the welfare-maximizing auctions by favoring an index representation of the optimal policy (more below). Similarly, the assumption that the value agent i derives from interacting with agent j is invariant in the composition of agent i 's matching set (that is, it does not depend on who else the individual interacts with) plays a similar role to the assumption that ε_{ij}^A and $\varepsilon_{ij'}^A$ are independent. It favors a description of the matching dynamics under profit-maximizing and welfare-maximizing auctions in terms of a simple index policy, when processes are endogenous. Such assumption can also be dispensed with in the case of exogenous processes, albeit at the cost of an increase in the complexity of the scores.²¹

The assumption that, in case of endogenous processes, kernels depend only the number of past interactions also plays the role of permitting a description of the matching dynamics in terms of (appropriately defined) indexes.

²⁰Specifically, each signal ξ_{ij}^A can be written as $\xi_{ij}^A = u_{ij}^A + \varsigma_{ij}^A$ with the innovations ς_{ij}^A drawn from a Normal distribution with mean 0 and variance ϑ_{ij}^A , independently from all other random variables.

²¹The case where each agent values interacting with at most one agent from the opposite side (one-to-one matching) corresponds to a particular relaxation of this assumption.

The assumption that the values agent i derives from interacting with agent j are independent of the values that other agents derive from interacting with the same agent in turn avoids the possibility that the platform trivially extracts the entire surplus from each agent using payments similar to those in Cremer and McLean (1988). Similarly, the assumption that agents possess some private information, θ_i^A , prior to entering into negotiations with the platform guarantees that the matching dynamics under profit maximization differ from their counterparts under welfare maximization. Together with the fact that each θ_i^A is unidimensional, this assumption permits us to shed light on distortions due to market power, while retaining tractability.

We see the combination of the above assumptions as a convenient tool to retain tractability in a fairly rich and complex environment while at the same time permitting us to shed light on the key trade-offs that platforms face in designing dynamic matching protocols when agents' true and/or perceived attractiveness changes over time.

Capacity constraints

In each period $t \geq 1$, the platform can match at most M agents from opposite sides, independently of past matching allocations. That is, M is a constraint on the number of matches the platform can entertain within each period. In each period, the platform can delete some of the previously formed matches and create new ones. We only impose that the total number of existing matches be no greater than M in all periods. As mentioned in the Introduction, such limited capacity may capture various constraints on the platform's side, such as limited facilities, time, or services availability.

We consider both the case in which M is sufficiently large that this constraint never binds (i.e., $M \geq n_A \cdot n_B$), as well as the case $M < n_A \cdot n_B$ in which this constraint may be binding. A period- t match $x_t \in \prod_{(i,j) \in N_A \times N_B} X_{ijt}$ is feasible if $\sum_{i \in N_A} \sum_{j \in N_B} x_{ijt} \leq M$. We denote by X_t the set of feasible period- t matches, and by $X \equiv \prod_{t=1}^{\infty} X_t$ the set of sequences of feasible matching allocations.

Matching mechanisms

The interactions the platform induces between the two sides can be thought of as governed by a *matching mechanism* $\Gamma \equiv (\mathcal{M}, \mathcal{S}, \chi, \psi, \rho)$. The latter consists of: (i) a collection of message sets $\mathcal{M} \equiv (\mathcal{M}_t)_{t=0}^{\infty}$, with each $\mathcal{M}_t \equiv \prod_{l \in N_k, k=A,B} \mathcal{M}_{lt}^k$ denoting the set of messages the agents can send in period t ; (ii) signals $\mathcal{S} \equiv (\mathcal{S}_t)_{t=0}^{\infty}$ that the platform may disclose to the agents, with $\mathcal{S}_t \equiv \prod_{l \in N_k, k=A,B} \mathcal{S}_{lt}^k$; (iii) a *matching rule* $\chi \equiv (\chi_t)_{t=1}^{\infty}$ describing, for each $t \geq 1$, the matches $\chi_t : \mathcal{M}^t \rightarrow X_t$ implemented given the history of received messages; (iv) a *payment rule* $\psi \equiv (\psi_t)_{t=0}^{\infty}$ describing, for each $t \geq 0$, the payments $\psi_t : \mathcal{M}^t \rightarrow \mathbb{R}^{N_A + N_B}$ asked from, or made to, the agents; and (v) a *disclosure policy* $\rho \equiv (\rho_t)_{t=0}^{\infty}$ specifying the information $\rho_t : \mathcal{M}^t \rightarrow \mathcal{S}_t$ disclosed to the agents over time. Each ρ_t must reveal to each agent his own matches. It may also reveal additional information, but it cannot conceal the matches in which the individual is involved.

A matching mechanism Γ is *feasible* if, for any sequence of messages $m \in \mathcal{M}$, the implemented allocations are feasible, that is, $\chi(m) \in X$.²²

²²Note that the rule χ is deterministic. This is because, in this environment, the platform never gains from inducing random matches, irrespective of whether its objective is profit maximization, or welfare maximization — see the discussion in the proof of Theorem 2 below.

Each agent chooses whether or not to join the platform after observing his vertical type θ_i^A but before observing his horizontal types. As anticipated above, this assumption reflects the idea that, in most markets of interest, agents choose whether or not to get “on-board” before learning the profile of potential partners. In environments in which the platform maximizes welfare (a possibility we consider below), the information the agents possess at the time they join the platform plays no role (they can be assumed to know only their vertical types θ , or also the period-1 horizontal types ε_1 , or to have no private information at all). If, instead, the platform maximizes profits, then the private information the agents possess at the time they join plays a fundamental role. If the agents possess no private information at all, the platform can extract the entire surplus. If, in addition to their vertical types, they also possess private information about their period-1 types ε_1 , the optimal mechanism is significantly more complicated, because of the multi-dimensionality of the agents’ initial private information (for example, it may be impossible to determine which participation constraints bind). To maintain tractability of the analysis, we thus assume that the agents’ private information at the time of joining is limited to the vertical types θ .

Upon joining the platform, at each period $t \geq 0$, each agent $l \in N_k$, from each side $k = A, B$, after learning $(\varepsilon_{ljt}^k)_{j \in N_{-k}}$, sends a message m_{lt}^k from the set \mathcal{M}_{lt}^k . In the auctions we introduce in the next section, such messages correspond to the selection of a membership status along with a collection of bids, one for each possible partner. Note that, while the game starts in period 0, the actual matching begins in period 1.

We use perfect Bayesian equilibrium (PBE) as the solution concept. A matching mechanism will be referred to as *profit-maximizing* if it is feasible and admits a PBE such that, under such equilibrium, the platform’s profits are at least as high as under any other PBE of any other feasible mechanism. A *welfare-maximizing* mechanism is defined in a similar way, but with welfare replacing profits in the platform’s objective.

Remark. The equilibria in the proposed matching auctions will satisfy properties stronger than those required by PBE. In particular, the equilibrium strategies remain optimal no matter the information each agent may possess about other agents’ past messages and match allocations, and no matter the beliefs the agent may have about other agents’ past and current types. Consistently with the rest of the dynamic mechanism design literature, we will refer to such equilibria as *periodic ex-post* (see, e.g., Bergemann and Valimaki (2010), Athey and Segal (2013), and Pavan, Segal, and Toikka (2014)). As a result, both the profit-maximizing and the welfare-maximizing mechanisms described below can be made *fully transparent*, in the sense that all bids, matches, and payments are made public after each period. Lastly, while we only require that the mechanisms induce participation in period zero, in the specific mechanisms we propose, agents find it optimal to remain on board after each history.

4 Matching auctions

We now introduce a class of matching mechanisms in which agents bid repeatedly to be matched with potential partners and where the matches with the highest (non-negative) bilateral scores are implemented in each period, up to capacity. The structure is the following.

Definition 1 (matching auctions) *In a matching auction:*

- At period $t = 0$, upon joining the platform, each agent $l \in N_k$, from each side $k = A, B$, is asked to select a membership status. Each status is conveniently indexed by the vertical type $\theta_{l0}^k \in \Theta_l^k$ it is designed for. A higher status translates into a more favorable treatment, on average, in the subsequent matching process. Accordingly, a higher status comes with a higher fee $p_{l0}^k \in \mathbb{R}$.

- In each subsequent period $t \geq 1$, each agent $l \in N_k$, from each side $k = A, B$, is first offered the possibility to revise his membership status by selecting $\theta_{lt}^k \in \Theta_l^k$. Each agent is then invited to submit bids $b_{lt}^k \equiv (b_{ijt}^k)_{j \in N_{-k}}$, one for each potential partner from the opposite side. Each pair of agents $(i, j) \in N_A \times N_B$ is then assigned a joint “score” $S_{ijt} \in \mathbb{R}$ that depends on the pair’s period-0 and period- t status, the pair’s period- t reciprocal bids (b_{ijt}^A, b_{ijt}^B) and, possibly, the number of times the pair interacted in the past. All pairs of agents with the highest nonnegative score are matched, up to capacity. All agents who are unmatched in period t , pay nothing. Matched agents make (or receive) payments $p_{it}^k \in \mathbb{R}$ that may depend also on other agents’ bids and membership status.

- All past bids, payments, membership choices, and matches are public.

The details of the scoring rules and of the corresponding payment functions naturally depend on the environment under examination and, in particular, the platform’s capacity constraint and the nature of the processes governing the evolution of the agents’ match values and the platform’s costs.

4.1 Scoring rules

A *scoring rule* S consists of a sequence of mappings $(S_{ijt})_{t=1, \dots, \infty}^{(i,j) \in N_A \times N_B}$ assigning to each pair of agents from opposite sides $(i, j) \in N_A \times N_B$, each $t \geq 1$, a score $S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) \in \mathbb{R}$ that depends on the pair’s period-0 and period- t membership status, current bids, and number of past interactions. Importantly, each score S_{ijt} depends only on information pertaining to the pair (i, j) . The conditioning of S_{ijt} on the entire profile $(\theta_0, \theta_t, b_t, x^{t-1})$ is only to facilitate the notation.

Formally speaking, the matching auctions defined above correspond to a particular class of matching mechanisms in which the message spaces are given by $\mathcal{M}_{i0}^A = \Theta_i^A$ for $t = 0$, and, for any $t \geq 1$, by $\mathcal{M}_{it}^A = \Theta_i^A \times \mathbb{R}^{N_B}$, and where the matching rule χ matches in each period $t \geq 1$ all pairs with the highest nonnegative score S_{ijt} , up to capacity, with ties broken arbitrarily. Obviously, the history of past matches $x^{t-1} = \chi^{t-1}(\theta^{t-1}, b^{t-1})$ is itself determined by past bids and membership choices. With an abuse of notation, whenever there is no risk of confusion, hereafter we drop the dependence of x^{t-1} on past bids and membership choices to ease the notation.

The matching rule χ corresponding to the scoring rule S then satisfies the following properties. If the pair (i, j) is matched in period t , it must be that its period- t score is nonnegative and that there are at most $M - 1$ matches with a period- t score strictly higher than the one for (i, j) . Furthermore, the pair (i, j) is *necessarily* matched in period t if its period- t score is strictly positive, and there are at most M matches with a period- t score weakly higher than that for (i, j) . Lastly, if there are no more than $M - 1$ matches with a score strictly higher than that of (i, j) , (i, j) ’s score is strictly positive, and the match (i, j) is not formed, then it must be that capacity is fully utilized and that the platform is matching other pairs with a score identical to (i, j) ’s.

Importantly, agents’ bids are allowed to be negative. An agent submitting a negative bid for a potential partner may nonetheless be matched with that partner if the platform finds it optimal to cross-subsidize the interaction. In considering the profitability of such cross-subsidization, the platform may also take into account

the value of experimenting by generating information about the profitability of such a match, which can be used in future periods.

We now introduce two scoring rules that play a central role in the analysis.

Definition 2 (myopic rule) *A myopic scoring rule $S^{m;\beta}$ (with weights β) is one in which, for each $t \geq 1$, each pair $(i, j) \in N_A \times N_B$, the score is given by*

$$S_{ijt}^{m;\beta} \equiv \beta_i^A(\theta_{i0}^A) \cdot b_{ijt}^A + \beta_j^B(\theta_{j0}^B) \cdot b_{ijt}^B - c_{ijt}(x^{t-1}), \quad (3)$$

where $\beta \equiv (\beta_l^k)_{l \in N_k, k=A,B}$ are time-invariant, non-decreasing, strictly positive, and bounded functions of the period-0 membership statuses. A myopic matching rule $\chi^{m;\beta}$ is a matching rule in which the scoring rule is given by $S^{m;\beta}$.

A myopic score is thus a weighted average of the bids b_{ijt}^A and b_{ijt}^B submitted by the pair $(i, j) \in N_A \times N_B$ in period t , $t \geq 1$, where the weights are given by strictly positive and non-decreasing functions of the agents' period-0 membership statuses, net of the platform's period- t cost of matching the pair. Note that myopic scores may depend on past bids only through the induced past allocations x^{t-1} , and are invariant in all membership choices except the period-0 ones. In the case of exogenous processes, because costs do not depend on past allocations, the myopic scores are also invariant in past bids.

The next scoring rule is a forward-looking one, where the scores take into account not only the flow profitability of the current interactions, but also the value of generating information that can be used in future periods. Let $\lambda_{ij}|\theta_t, b_t, x^{t-1}$ denote the stochastic process over the pair (i, j) 's current and future match values (v_{ijs}^A, v_{ijs}^B) and costs c_{ijs} that one obtains under any matching rule that matches the pair (i, j) at all periods $s \geq t$, when the pair's vertical types (θ_i^A, θ_j^B) are the ones revealed by the period- t membership choices θ_t , and when the horizontal types are given by

$$\varepsilon_{ijt}^k = \begin{cases} b_{ijt}^k / \theta_{it}^k & \text{if } b_{ijt}^k / \theta_{it}^k \in \mathcal{E}_{ijt}^k \\ \arg \min_{\varepsilon_{ijt}^k \in \mathcal{E}_{ijt}^k} \left\{ |b_{ijt}^k / \theta_{it}^k - \varepsilon_{ijt}^k| \right\} & \text{otherwise,} \end{cases} \quad (4)$$

$k = A, B$.

Definition 3 (index rule) *An index scoring rule $S^{I;\beta}$ (with weights β) is one in which, for each $t \geq 1$, each pair $(i, j) \in N_A \times N_B$,*

$$S_{ijt}^{I;\beta} \equiv \max_{\tau} \left\{ \frac{\mathbb{E}^{\lambda_{ij}|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\tau} \delta^{s-t} \left(\beta_i^A(\theta_{i0}^A) v_{ijs}^A + \beta_j^B(\theta_{j0}^B) v_{ijs}^B - c_{ijs} \right) \right]}{\mathbb{E}^{\lambda_{ij}|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\tau} \delta^{s-t} \right]} \right\}, \quad (5)$$

where τ denotes a stopping time, and where $\beta \equiv (\beta_l^k)_{l \in N_k, k=A,B}$ are time-invariant, non-decreasing, strictly positive, and bounded functions of the period-0 membership statuses. An index matching rule $\chi^{I;\beta}$ is a matching rule in which the scoring rule is given by $S^{I;\beta}$.

The period- t score $S_{ijt}^{I;\beta}$ for the match between agent i from side A and agent j from side B thus corresponds

to a Gittins index for a process for which the “rewards” are given by the weighted total surplus of the match (i, j) , with the weights given by the functions β of the agents’ period-0 membership statuses. Note that, under our assumptions for endogenous processes, each score $S_{ijt}^{I;\beta}$ depends on x^{t-1} only through the total number of past interactions $\sum_{s=1}^{t-1} x_{ijs}$ between the pair (i, j) . Also note that, contrary to the myopic scores, the indexes S_{ijt}^I depend on the entire history of past and current membership choices. However, while the dependence on the period-0 and on the period- t choice is direct, the dependence on other periods’ choices is only through the number of past interactions.

4.2 Role of the membership status

Intuitively, each agent’s membership status determines the importance the platform assigns to the agent’s bids, relative to those of others. Consider, for example, an environment with exogenous processes, and suppose that matches are determined according to the myopic rule $\chi^{m;\beta}$. Suppose, in period $t \geq 1$, agent i from side A submits a positive bid for agent j from side B , whereas agent j submits a negative bid. For given bids (b_{ijt}^A, b_{ijt}^B) , a higher period-0 status of agent i implies a higher score for the match (i, j) , thus tilting the matching allocation in favor of agent i . Symmetrically, a higher status for agent j reduces the score for the match (i, j) , thus tilting the allocation in favor of agent j . A higher period-0 status thus grants an agent preferential treatment in all subsequent auctions, both with respect to the competition the agent faces with other agents from his own side (when the capacity constraint is binding) and with respect to the competition the agent faces with agents from the opposite side, shifting the length of each interaction to his benefit.

The role of the agents’ status in subsequent periods (that is, θ_{it}^A , with $t \geq 1$), is somewhat different. It permits the platform to decouple the agents’ horizontal types ε_{ijt}^A from the agents’ bids, which is useful to form expectations about the agents’ future match values. In fact, as we show below, while the agents’ bids convey information about the agents’ total match values v_{ijt}^A , they are not sufficient statistics with respect to the agents’ overall private information, when it comes to predicting future values. For instance, a high bid b_{ijt}^A by agent i for agent j may either reflect a persistent high value for interacting with all agents from the opposite side (i.e., a high vertical type θ_i^A), or a high temporary appreciation for interacting with agent j (i.e., a high horizontal type, ε_{ijt}^A). Accordingly, when the evolution of the agents’ match values is endogenous, the platform uses the combination of the submitted bids with the membership statuses to trade off the value of current matches with the value of generating information useful in future periods.

Finally, note that the reason why the platform allows the agents to revise their status over time, despite the fact that, under the specification in (1), the vertical types are perfectly persistent, is that this favors equilibria in which the agents bid myopically over time, irrespective of past bids and membership selections, in the following sense:

Definition 4 (truthful bidding) *A strategy profile $\sigma = (\sigma_l^k)_{l \in N_k}^{k=A,B}$ for the above matching auctions is truthful if for all $t \geq 0$, all private histories, the selection of the membership status is given by $\theta_{it}^k = \theta_l^k$, and the submitted bids are given by $b_{ijt}^k = v_{ijt}^k = \theta_l^k \cdot \varepsilon_{ijt}^k$, all $(i, j) \in N_A \times N_B$, $k = A, B$. A truthful equilibrium is an equilibrium in which the strategy profile is truthful.*

That is, under truthful strategies, irrespective of the history, agents select a membership status equal to

their true vertical type, and then submit bids equal to the myopic match values they assign to each interaction.

4.3 Payments

We now complete the description of the matching auctions by describing the associated payment rules. Let χ be the matching rule corresponding to the scoring rule S . For any $t \geq 1$, any $(\theta_0, \theta_t, b_t, x^{t-1})$, any weights $\beta \equiv (\beta_l^k)_{l \in N_k, k=A,B}$, let

$$W_t \equiv \mathbb{E}^{\lambda[\chi]|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i \in N_A} \sum_{j \in N_B} S_{ijs}^{m;\beta} \cdot \chi_{ijs} \right] \quad (6)$$

denote the continuation weighted surplus, that is, the net present value of current and future myopic scores, implemented under the rule χ . Here $\lambda[\chi]|\theta_t, b_t, x^{t-1}$ denotes the stochastic process over (θ_s, b_s, x^{s-1}) , $s \geq t$, under the rule χ , when the selected period- t membership statuses are θ_t , the period- t bids are b_t , all agents follow truthful strategies from period $s > t$ onwards, the true vertical types are the ones corresponding to the selected period- t statuses (i.e., $\theta_l^k = \theta_{lt}^k$, $l \in N_k$, $k = A, B$), and the true period- t horizontal types are given by (4).

Similarly, let $W_t^{-l,k}$ denote the continuation weighted surplus, as defined in (6), when the myopic scores involving agent l from side k are identically equal to zero at all periods $s \geq t$.

Next, let $R_{lt}^k \equiv W_t - W_t^{-l,k}$, denote the *contribution* of agent $l \in N_k$ to the continuation weighted surplus and

$$r_{lt}^k \equiv R_{lt}^k - \delta \mathbb{E}^{\lambda[\chi]|\theta_t, b_t, x^{t-1}} \left[R_{lt+1}^k \right] \quad (7)$$

the corresponding *flow marginal contribution*.

In each period $t \geq 1$, the payment asked to each agent $i \in N_A$ (an analogous construction holds for each agent from side B) is given by

$$\psi_{it}^A = \sum_{j \in N_B} b_{ijt}^A \cdot \chi_{ijt} - \frac{1}{\beta_i^A(\theta_{i0}^A)} r_{it}^A. \quad (8)$$

In words, each agent i is asked to make a payment that equals the total flow value the agent derives from all the matches implemented in period t in which the agent is involved (as reflected by the agent's own bids), net of a discount that is proportional to the agent's flow marginal contribution to weighted surplus, with a coefficient of proportionality given by $1/\beta_i^A(\theta_{i0}^A)$.

The description of the payments is completed by the specification of the fees charged to the agents in period zero, as a function of the profile of selected membership statuses. The latter are given by

$$\psi_{i0}^k = \theta_{i0}^k D_i^k(\theta_0) - \int_{\theta_i^k}^{\theta_{i0}^k} D_i^k(\theta_{-l}^k, y) dy - \mathbb{E}^{\lambda[\chi]|\theta_0} \left[\sum_{t=1}^{\infty} \delta^t \psi_{it}^k \right] - L_l^k, \quad (9)$$

where L_l^k is a scalar whose role is to guarantee that it is optimal for the agent to participate, θ_{-l}^k is a profile of membership statuses for all agents excluding agent l from side k , and

$$D_i^k(\theta_0) \equiv \begin{cases} \mathbb{E}^{\lambda[\chi]|\theta_0} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N_{-k}} \varepsilon_{lht}^k \chi_{lht} \right] & \text{if } k = A \\ \mathbb{E}^{\lambda[\chi]|\theta_0} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N_{-k}} \varepsilon_{hlt}^k \chi_{hlt} \right] & \text{if } k = B \end{cases} \quad (10)$$

is the “quality” of the matches agent l from side k expects from joining the platform, when the profile of selected period-0 membership statuses is θ_0 (here $\lambda[\chi]|\theta_0$ denotes the distribution over future membership statuses, bids, and matches, under truthful strategies). Note that the reason we distinguish between the case in which $k = A$ and the one in which $k = B$ is that the order in the subscripts of the allocations χ_{ijt} , as well as the order in the subscripts in the horizontal types ε_{ijt}^k is not permutable: the first index always refers to side A , while the second to side B . Also note that, according to (9), the fee the platform charges to each agent depends on the entire profile of membership statuses selected by all agents. In other words, the price for status depends on the aggregate demand of status from each side of the market. Such dependence plays a role analogous to that of “insulating tariffs” in two-sided markets (see, e.g., Weyl, 2010): it guarantees that all agents find it optimal to join the platform in period zero, irrespective of their beliefs about the types of other agents. Such insulating tariffs can always be replaced by (perhaps more familiar) tariffs in which membership fees depend only on each agent’s own status, by taking expectations over the types of other agents.

The period- t payments, $t \geq 1$, on the other hand, reflect the externalities the agents impose both in the current and in future periods on all other agents (from both sides of the market), as well as the costs they impose on the platform. Such externalities are calculated with respect to the “weighted surplus” associated with each of the matches, with the weights determined by the period-0 membership statuses. Note that such externalities may be positive or negative, and therefore the payments may be positive or negative. For example, if an agent is valued highly by the agents he is matched to, he may receive a positive transfer from the platform – cross subsidization.

As for the membership fees, because higher status implies higher average match quality (a property we formally establish in the proof of Theorem 1 in the Appendix), higher status naturally comes with a higher fee. In particular, the period-0 membership fees are equal to the agents’ expected match quality, net of the payments the agents expect to make in the subsequent periods, and net of discounts $\int_{\theta_l^k}^{\theta_{l0}^k} D_l^k(\theta_{-l}^k, y) dy$ whose role is to incentivize the agents to select the status corresponding to their true vertical types.

Interestingly, contrary to the payments in standard Vickrey-type auctions such as the GSP auctions used for sponsored search, multiple agents (from both sides of the market) may be charged for the same externality they impose on others.

5 Equilibrium

We now turn to the equilibria of the matching auctions introduced above. We start by describing environments that satisfy the separability condition informally introduced in the Example in Section 2. Let $S^{I;\beta}$ denote the index scoring rule defined by the weights β (as formally introduced in Definition 3). Then let

$$\underline{S}_{ijt}^\beta \equiv \inf_{s \leq t} \{ S_{ijs}^{I;\beta} \} \quad (11)$$

denote the historical infimum of the indexes for the pair (i, j) .²³

²³Such auxiliary processes are also referred to as *lower envelope processes* (see, e.g., Mandelbaum 1986).

Definition 5 (separability) *The environment is separable under the rule $S^{I;\beta}$ if, for any $t \geq 1$, any (θ^t, b^t, x^{t-1}) , any two pairs of agents $(i, j), (i', j') \in N_A \times N_B$ for which $\underline{S}_{ijt}^\beta$ and $\underline{S}_{i'j't}^\beta$ are, respectively, the M -th and $(M+1)$ -th highest are such that*

$$\underline{S}_{ijt}^\beta > \underline{S}_{i'j't}^\beta > 0 \Rightarrow \underline{S}_{ijt}^\beta \cdot (1 - \delta) \geq \underline{S}_{i'j't}^\beta. \quad (12)$$

Note that the definition imposes a restriction on the initial heterogeneity of the positive indexes, as well as on the way such indexes decrease over time. It imposes no restriction on the way the indexes respond to positive shocks, or on the way they evolve once they turn negative. The following example of a separable environment generalizes the one in Section 2.

Example 3 (bad news) Consider an environment in which the period-1 indexes $S_{ij1}^{I;\beta}$ are sufficiently heterogeneous, in the sense of (12). Further assume that, for each pair of agents (i, j) , in each period $t \geq 2$, either $S_{ijt}^{I;\beta} \geq S_{ijt-1}^{I;\beta}$ or $S_{ijt}^{I;\beta} < 0$. The environment is then separable. The above structure emerges, for example, in the familiar “no-news-is-good-news” class of environments considered in the experimentation literature (e.g., Bonatti and Horner (2015), Keller and Rady (2015)).²⁴ In such environments, in each period, agents either receive no news or negative news revealing the unprofitability of the matches. In the absence of bad news, agents become increasingly more optimistic about their relationships.²⁵ \square

We then have the following result:

Theorem 1 (equilibrium) *(i) Suppose the processes are exogenous. Any matching auction in which (1) the scoring rule is myopic with arbitrary weights β and (2) the payments are given by (8) and (9), with the same weights β as in the scoring rule and with L_l^k large enough, all $l \in N_k, k = A, B$, admits an equilibrium in which all agents participate in each period and follow truthful strategies.*

(ii) Suppose processes are endogenous and either (a) $M = 1$, or (b) $M \geq n_A \cdot n_B$, or (c) $1 < M < n_A \cdot n_B$ and, in this latter case, the environment is separable under the rule $S^{I;\beta}$. Any matching auction in which (1) the scoring rule is the index rule with weights β and (2) the payments are given by (8) and (9), with the same weights β as in the scoring rule, and with L_l^k large enough, all $l \in N_k, k = A, B$, admits an equilibrium in which all agents participate in each period and follow truthful strategies.

(iii) The equilibria of the auctions in parts (i) and (ii) above are periodic ex-post; that is, the agents’ strategies are sequentially rational, regardless of the agents’ beliefs about other agents’ past and current types.

Heuristic Proof. The formal proof is in the Appendix. Here we illustrate the key ideas in an heuristic way. We organize the arguments in three steps. The first step shows that, in the continuation game that starts with period $t \geq 1$, independently of past behavior, when the agents follow truthful strategies, the matches implemented under the myopic and the index rules maximize continuation weighted surplus, as defined in (6), in the respective environments of Theorem 1. Step 2 in turn establishes that, when the payments are the ones

²⁴See Keller, Rady and Cripps (2005) for the opposite case of fully revealing “breakthroughs” and Horner and Skrzypacz (2016) for a recent survey of the literature on experimentation in strategic environments.

²⁵A similar structure obtains in environments in which, in each period and for each match, with some positive probability (possibly history dependent), the match is revealed to be unprofitable, in which case the index turns negative, while with the complementary probability the index either increases or decreases by a small amount that does not violate (12).

in (8) above, in the continuation game that starts with period $t \geq 1$, irrespective of past behavior, all agents find it optimal to select a membership status equal to their true vertical type and bid truthfully their myopic values v_{ijt}^k . Step 3 completes the proof by showing that, when the period-0 membership fees are the ones in (9), all agents find it optimal to join the platform in period zero and select the period-0 membership status corresponding to their true vertical type, again, regardless of their beliefs about other agents' types.

Step 1. Suppose the processes are exogenous. That the matches implemented under truthful strategies maximize continuation weighted surplus follows from the fact that a myopic rule (with weights β) selects, after each history, the matches with the highest nonnegative myopic score.

The result for endogenous processes is more delicate. In this case, the problem of maximizing continuation weighted surplus is a *multiarmed bandit problem*. Such a problem admits an index-policy solution when either the capacity constraint is not binding ($M \geq n_A \cdot n_B$), or when the platform can match at most one pair of agents in each period ($M = 1$). In both such cases, continuation weighted surplus is maximized by selecting in each period the matches with the highest nonnegative index, up to capacity.

For intermediate capacity constraints ($1 < M < n_A \cdot n_B$), however, without further restrictions on the environment, there is no guarantee that an index policy maximizes continuation weighted surplus. This is because, as anticipated in the Introduction and further discussed in Section 2, *in general*, multiarmed bandit problems in which multiple arms can be activated simultaneously fail to admit a simple index solution. The Supplementary Material contains an example of an environment in which the separability condition of Definition 5 is violated and illustrates the suboptimality of the index policy in the context of our matching environment. In the case of intermediate capacity constraints, the result in Theorem 1 above thus requires the environment to be separable under the rule $S^{I;\beta}$, in the sense of Definition 5. As we show in the Appendix, the role of the separability condition is to guarantee that the optimal rule is myopic in a fictitious environment in which the reward processes are the auxiliary ones, $\underline{S}_{ijt}^\beta$, as defined in (11). The proof then proceeds by showing that, given any matching rule χ , continuation weighted surplus is weakly higher when the rewards are given by the auxiliary processes (i.e., $\underline{S}_{ijt}^{m;\beta}$) than when the rewards are the primitive ones (i.e., $S_{ijt}^{m;\beta}$). Furthermore, in the special case in which the scoring rule is the index rule (for the primitive environment), continuation weighted surplus is the same in the two environments. We then show that when the environment is separable, a myopic rule maximizes continuation weighted surplus when the rewards are given by the auxiliary processes. Recall that such processes drift downwards. The maximal continuation surplus that can be expected from each match is thus equal to the “annuity” $\underline{S}_{ijt}^\beta / (1 - \delta)$ of its current auxiliary value. Separability then implies that it is always optimal to select the matches for which the current auxiliary reward is the highest.

That the matches implemented by the index rule maximize continuation weighted surplus in the primitive environment then follows from the above properties along with the fact that the matches implemented by such rule in the primitive environment coincide with those implemented by the myopic rule in the fictitious environment.

Step 2. The second step of the proof shows that, when the payments are the ones specified in (8), starting from each period $t \geq 1$, and irrespective of past behavior, all agents have incentives to select a membership status equal to their true vertical type and then bid truthfully their myopic values $b_{ijt}^k = v_{ijt}^k$ in the continuation

game that starts with period t . This step combines arguments from Bergemann and Valimaki (2010) (see also Kakade et al. (2013)) with arguments from Pavan, Segal, and Toikka (2014). A notorious difficulty in dynamic incentives problems is the need to control for multi-period contingent deviations. In matching auctions, agents who deviated from a truthful strategy in the past may wish to do so again at present and/or in future periods. An additional difficulty in the matching environment under consideration here is the multi-dimensionality of the new private information each agent receives in each period.

Given the profile of period-0 membership choices, the payments in (8) are designed to make each agent's continuation payoff (net of the payments) proportional to his flow marginal contribution to *weighted surplus* (with the latter defined as in (7), and with the coefficient of proportionality given by $1/\beta_t^k(\theta_{i0}^k)$). Because the matches under the proposed scoring rules maximize the continuation weighted surplus after all histories (including those off the equilibrium path), agents have incentives to go back to the truthful strategies (and remain in the mechanism) after all histories, and irrespective of the beliefs they may have about past and current types of all other agents. In other words, participating and then following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any history. Importantly, when processes are endogenous, that agents have the possibility to adjust their membership status at the beginning of each period helps establishing the sequential optimality of truthful strategies by guaranteeing that, irrespective of past behavior, the matches implemented in the continuation game under truthful strategies maximize the continuation weighed surplus, at all histories.²⁶

Step 3. The final step completes the proof by showing that, when the membership fees are the ones in (9), all agents find it optimal to join the platform in period zero, and select the membership status designed for their true vertical type, regardless of the agents' beliefs about other agents' types.

Specifically, the period-0 payments are added to the payments in the subsequent periods so that the payoff that each agent obtains in equilibrium satisfies a certain envelope condition, which relates the agent's period-0 interim expected payoff to the expected discounted "match quality," as defined in (10). The proof then shows that the myopic and index rules satisfy an "average monotonicity" property, according to which match quality increases both with each agent's period-0 membership status, for fixed true vertical type, and with the agent's true vertical type, under truthful strategies. Along with the fact that the equilibrium payoffs satisfy the aforementioned envelope formula, such "average monotonicity" property guarantees that truthful strategies are optimal also in period zero.

Importantly, the quality of an agent's interactions need not increase with his membership status, or with his true vertical type, in each state of the world (if fact, it does not, under the index rule²⁷). It suffices that it increases, on average, where the averaging is across time and states.²⁸ ■

²⁶A similar approach of enlarging the message space so as to give agents the possibility to reveal their true types after possible deviations in previous periods is used in Doepke and Townsend (2006) and Kakade et al. (2013).

²⁷This is because an increase, for example, in agent i 's period-0 membership status may lead to reversals in the ordering of the agent's indexes. Because these indexes are forward looking, such reversals, while optimal based on the period- t information, might not be optimal ex-post. As a result, the monotonicity of an agent's discounted sum of current and future horizontal types under an index rule need not hold ex-post. That is, along certain paths, higher membership status may result in lower match quality.

²⁸The above condition can be viewed as the analog of Pavan, Segal and Toikka (2014)'s average monotonicity condition in a setting, the present one, in which agents receive multi-dimensional new private information in each period. In fact, in the present environment, the horizontal types $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N_B}$ correspond to the "impulse responses" of the period- t match values $v_{it}^A \equiv (v_{ijt}^A)_{j \in N_B}$ to the vertical type θ_i^A , as defined in Pavan, Segal, and Toikka (2014).

In environments in which agents have a fixed demand for interactions with each potential partner (as in Example 2 above), the separability condition in part (ii) of the theorem can be relaxed. More generally, separability is sufficient for the index rule to maximize continuation weighted surplus, and hence for the agents to find it optimal to follow truthful strategies in the proposed matching auctions, but it is not necessary.²⁹

Also, note that the periodic ex-post nature of the equilibria of Theorem 1 is what justifies making the auctions fully transparent.

A key difficulty when agents learn their match values by interacting with other agents is that the private value of experimentation need not align across partners (a problem that does not emerge in standard auctions for physical goods). For example, after a few interactions, agent i from side A may have learned his value for agent j from side B , while agent j may face residual uncertainty about his value for agent i . The scores $S^{I;\beta}$ in the proposed auctions are thus different from a simple combination of the Gittins indexes corresponding to the agents' own private value for experimentation. They are constructed to internalize the cost and benefits of cross-subsidization, as perceived by the platform, while also accounting for the costs of the agents' private information.

One of the appeals of the proposed auctions is that they admit simple equilibria in which agents bid their myopic match values in each period. In other words, the agents do not need to know how to solve complex dynamic-programming problems, or be able to compute the indexes (although, nowadays, there is software that does so). Once the scoring and payments rules are understood, it is in the agents' interest to bid "straightforwardly" in all periods, irrespective of the beliefs they may have about other agents' past and current types, and irrespective of their own, as well as other agents', past behavior. As in other market design settings, it is, however, important that the market designer educates the bidders by carefully explaining the structure of the scoring and payment rules and why the proposed auctions admit such simple equilibria. Also note that the proposed auctions are, in spirit, the analogs of the familiar second-price Vickrey auctions, but adapted to control for the dynamic nature of the externalities the agents impose on one another, and for the fact that the matches implemented in equilibrium need not maximize total surplus.

6 Profit maximization

We now show that the matching auctions introduced above include a subclass that maximizes the platform's profits across all possible mechanisms. We make the additional assumption that, for all $l \in N_k$, $k = A, B$, the Mills ratio $[1 - F_l^k(\theta_l^k)] / f_l^k(\theta_l^k)$ is non-increasing and that $\theta_l^k \cdot f_l^k(\theta_l^k) > 1$.³⁰

Because the myopic and the index rules are completely characterized by the weights β , hereafter, we denote by $D_l^k(\theta_0; m, \beta)$ and by $D_l^k(\theta_0; I, \beta)$, respectively, the expected match quality under a myopic and under an index rule, with weights β . We then have the following result:

²⁹To the best of our knowledge, a condition that is jointly necessary and sufficient for the optimality of index policies in multi-armed bandit problems in which an arbitrary number of arms is activated in each period remains elusive, except in special environments.

³⁰The first part of the assumption is standard in mechanism design and guarantees that the agents' "virtual" vertical types $\theta_l^k - [1 - F_l^k(\theta_l^k)] / f_l^k(\theta_l^k)$ are non-decreasing in the true types. The second part is added to guarantee that the virtual vertical types are strictly positive. Given the multiplicative structure of the match values in (1), this assumption guarantees that the virtual match values $[\theta_l^k - [1 - F_l^k(\theta_l^k)] / f_l^k(\theta_l^k)] \varepsilon_{ijt}^k$ respect the same ranking as the true ones, $v_{ijt}^k = \theta_l^k \varepsilon_{ijt}^k$, thus avoiding confusion in the interpretation.

Theorem 2 (profit maximization) Let $\hat{\beta}$ be the weights given by³¹

$$\hat{\beta}_l^k(\theta_{l_0}^k) \equiv 1 - \frac{1 - F_l^k(\theta_{l_0}^k)}{f_l^k(\theta_{l_0}^k)\theta_{l_0}^k}, \text{ all } l \in N_k, k = A, B. \quad (13)$$

(i) Suppose processes are exogenous and $D_l^k(\theta_{-l}^k, \theta_l^k; m, \hat{\beta}) \geq 0$, all $l \in N_k$, $k = A, B$, and all θ_{-l}^k . The matching auctions in which (1) the scoring rule is the myopic rule with weights $\hat{\beta}$, and (2) the payments are given by (8) and (9) with weights $\hat{\beta}$ and with $L_l^k = 0$, all $l \in N_k$, $k = A, B$, are profit-maximizing.

(ii) Suppose processes are endogenous and $D_l^k(\theta_{-l}^k, \theta_l^k; I, \hat{\beta}) \geq 0$, all $l \in N_k$, $k = A, B$, and all θ_{-l}^k . In each of the three environments of part (ii) of Theorem 1, the matching auctions in which (1) the scoring rule is the index rule with weights $\hat{\beta}$, and (2) the payments are given by (8) and (9) with weights $\hat{\beta}$ and with $L_l^k = 0$, all $l \in N_k$, $k = A, B$, are profit-maximizing.³²

(iii) The equilibria of the matching auctions in parts (i) and (ii) above that maximize the platform's profits are in truthful strategies, and are periodic ex-post.

The proof of Theorem 2, in the Appendix, is in three steps. First, we show that, given any matching mechanism Γ and any *Bayes Nash equilibrium* (BNE) σ of the game induced by Γ , the period-0 interim expected payoff of each agent $l \in N_k$ from each side $k = A, B$ must be given by³³

$$U_l^k(\theta_l^k) = U_l^k(\theta_l^k) + \int_{\theta_l^k}^{\theta_l^k} \mathbb{E} \left[D_l^k(\theta; \hat{\chi}) | y \right] dy, \quad (14)$$

where the rule $\hat{\chi} = (\hat{\chi}_t(\theta, \varepsilon^t))_{t=1}^{\infty}$ describes the state-contingent matches induced by the strategy profile σ in Γ , and where the expectation in (14) is with respect to the entire profile of vertical types θ , given the agent's own vertical type.³⁴

Next, we use the above representation of the agents' equilibrium interim expected payoffs to show that, given any mechanism Γ and any BNE σ of Γ , the platform's profits are given by the following weighted surplus function

$$\begin{aligned} & \mathbb{E}^{\lambda[\hat{\chi}]} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i \in N_A} \sum_{j \in N_B} \left(\hat{\beta}_i^A(\theta_i^A) \theta_i^A \varepsilon_{ijt}^A + \hat{\beta}_j^B(\theta_j^B) \theta_j^B \varepsilon_{ijt}^B - c_{ijt}(\hat{\chi}_{ij}^{t-1}(\theta, \varepsilon^{t-1})) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right) \right] \\ & - \sum_{k=A,B} \sum_{l \in N_k} U_l^k(\theta_l^k), \end{aligned} \quad (15)$$

for which the first term is only a function of the matching rule, and with the weights $\hat{\beta}$ as given in (13).

³¹Note that, under the maintained assumption on the F distributions, these weights are non-decreasing, strictly positive, and bounded, as required by the scoring rules of Theorem 1.

³²If $1 < M < n_A \cdot n_B$, the optimality of the proposed matching auctions requires the environment to be separable with respect to the rule $S^{I; \hat{\beta}}$ corresponding to the weights $\hat{\beta}$.

³³By period-0 interim expected payoff, we mean the payoff the agent expects, under the equilibrium strategy profile σ , in the game induced by Γ , when his period-0 vertical type is equal to θ_l^k .

³⁴In case of endogenous processes, these functions are defined only for those histories that are consistent with equilibrium play in previous periods. To ease the exposition, hereafter we omit the formal specification of the domains of such functions.

The optimality of the auctions in Theorem 2 is then established by showing that, under the truthful equilibria of the proposed auctions, (a) the induced state-contingent matches maximize the first component of the weighted surplus function (15), and (b) the participation constraints for the lowest vertical types are binding. That the lowest vertical type of each agent expects a nonnegative match quality (formally, that $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; r, \hat{\beta}) \geq 0$, $r = m, I$), together with the fact that match quality $D_l^k(\theta_{-l}^k, \cdot; r, \hat{\beta})$, $r = m, I$, is non-decreasing in the agent's vertical type θ_l^k under the equilibria of Theorem 2 and that the interim expected payoffs satisfy the envelope conditions in (14), in turn guarantees that participation in period zero is optimal for all agents. In fact, we show that payoffs satisfy a stronger envelope condition, which guarantees the optimality of period zero participation irrespective of the agents' beliefs about other agents' types.

The conditions in Theorem 2 pertaining to expected match quality D_l^k can be relaxed for standard PBE solution concepts, and are vacuously satisfied, for example, if horizontal types are nonnegative, that is, if no agent ever dislikes interacting with any other agent. Interestingly, the truthful equilibria of the proposed auctions (with weights $\hat{\beta}$) maximize profits over all *BNE* (not just PBE) of all indirect mechanisms.

7 Welfare vs profit maximization

We now turn to the distortions in the dynamics of the matching allocations that emerge when the platform maximizes profits instead of welfare. Theorem 3 below identifies conditions under which the auctions introduced in Section 4 above admit a subclass for which the matches sustained in equilibrium are welfare maximizing.

Theorem 3 (welfare maximization) *Let β^W be the weights given by $\beta_l^{k,W}(\theta_l^k) = 1$, all θ_l^k , $l \in N_k$, $k = A, B$.*

(i) *Suppose processes are exogenous. The matching auctions in which (1) the scoring rule is the myopic rule with weights β^W , and (2) the payments are given by (8) and (9) with weights β^W and with L_l^k large enough, all $l \in N_k$, $k = A, B$, are welfare-maximizing.*

(ii) *Suppose processes are endogenous. In each of the three environments of part (ii) of Theorem 1, the matching auctions in which (1) the scoring rule is the index rule with weights β^W , and (2) the payments are given by (8) and (9) with weights β^W and with L_l^k large enough, all $l \in N_k$, $k = A, B$, are welfare-maximizing.³⁵*

(iii) *Suppose (a) processes are exogenous and $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; m, \beta^W) \geq 0$, all $l \in N_k$, $k = A, B$, and all θ_{-l}^k , or (b) processes are endogenous and $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; I, \beta^W) \geq 0$, all $l \in N_k$, $k = A, B$, and all θ_{-l}^k . In their respective environments of parts (i) and (ii) above, the matching auctions with payments given by (8) and (9), with $L_l^k = 0$, all $l \in N_k$, $k = A, B$, admit ex-post periodic equilibria in which agents participate and follow truthful strategies at all histories. Furthermore, such auctions maximize the platform's profits over all mechanisms implementing welfare-maximizing matches and inducing the agents to join the platform in period zero.*

The results in parts (i) and (ii) follow directly from Theorem 1 by noting that, when the weights are given by β^W , the matching allocations sustained under truthful strategies maximize welfare after each history. The conditions in part (iii) of Theorem 3 in turn guarantee that, when the payments are as in (8) and (9) with weights β^W and with $L_l^k = 0$, all $l \in N_k$, $k = A, B$, agents find it optimal to participate and follow

³⁵If $1 < M < n_A \cdot n_B$, the separability condition must hold with respect to the rule $S^{I;\beta^W}$ with weights β^W .

truthful strategies in each period, irrespective of their beliefs about other agents' current and past types, and past matches. That the proposed auctions maximize the platform's profits among all mechanisms implementing welfare-maximizing matches and inducing the agents to participate in period zero follows from arguments similar to those establishing the optimality of the matching auctions in Theorem 2 above. Note that the conditions in part (iii), though, do not guarantee that the platform's profits be positive. In general, the government may thus need to subsidize the platform to induce it to implement welfare-maximizing matching allocations, while also recouping its costs.

We now turn to the distortions brought in by profit maximization. Let $\chi^P \equiv (\chi_t^P(\theta, \omega))_{t=1}^\infty$ and $\chi^W \equiv (\chi_t^W(\theta, \omega))_{t=1}^\infty$ denote the state-contingent profit- and welfare-maximizing matching allocations, under the equilibria of the auctions of Theorems 2 and 3, respectively. Because the period- t history of horizontal types ε^t may differ in the two auctions (under endogenous processes), these allocations are expressed as a function of the sequence $\omega \equiv (\omega_{ijs}^k)_{(i,j) \in N_A \times N_B, k \in \{A, B\}, s=1, \dots, \infty}$ of exogenous random variables that, together with the history of past matches (in case of endogenous processes), generate the sequence of horizontal types ε .³⁶ Because (θ, ω) are exogenous, we then drop them from the arguments of χ^P and χ^W to ease the notation. We then have the following result:³⁷

Theorem 4 (distortions) *Suppose all agents derive a nonnegative utility from interacting with all other agents from the opposite side (formally, $\varepsilon_{ijt}^k \geq 0$, all $(i, j) \in N_A \times N_B$, $k = A, B$, $t \geq 1$).*

(1) *If $M \geq n_A \cdot n_B$, then under both exogenous and endogenous processes, for all $(i, j) \in N_A \times N_B$, all $t \geq 1$: $\chi_{ijt}^P = 1 \Rightarrow \chi_{ijt}^W = 1$.*

(2) *Suppose that either (a) the processes are exogenous and $M < n_A \cdot n_B$, or (b) the processes are endogenous and $M = 1$. Then, for all $t \geq 1$: $\sum_{(i,j) \in N_A \times N_B} \chi_{ijt}^W \geq \sum_{(i,j) \in N_A \times N_B} \chi_{ijt}^P$.*

(3) *Suppose processes are endogenous, the environment is separable under both $S^{I;\hat{\beta}}$ and $S^{I;\beta^W}$, and $1 < M < n_A \cdot n_B$. If, under profit-maximization, matching stops in finite time, then:*

$$\sum_{t=1}^{\infty} \sum_{(i,j) \in N_A \times N_B} \chi_{ijt}^W \geq \sum_{t=1}^{\infty} \sum_{(i,j) \in N_A \times N_B} \chi_{ijt}^P.$$

As in other screening problems, distortions are introduced under profit maximization to reduce the agents' information rents (that is, the surplus the platform must leave to the agents to induce them to reveal their private information). When all agents value positively interacting with all other agents from the opposite side, a profit-maximizing platform induces fewer interactions than what is efficient. The precise nature of the inefficiency, however, depends on whether the processes governing the evolution of the agents' match values are endogenous or exogenous, and on the platform's capacity constraint.

When the capacity constraint is not binding, irrespective of the nature of the processes, in each period $t \geq 1$, the set of matches accommodated by a profit-maximizing platform is a subset of those maximizing

³⁶In case of exogenous processes, one can think of the horizontal types ε as coinciding with the random variables ω . With endogenous processes, instead, the horizontal types ε are generated from the exogenous random variables ω and the past decisions using the construction explained in Section 3.

³⁷For this result, to facilitate the comparison between profit and welfare maximization, we assume that, in each period, both the profit- and the welfare-maximizing rules match pairs for which the scores are zero. That is, both rules use all the available M slots, unless the number of matches for which the score is nonnegative is strictly less than M .

welfare. In case of endogenous processes, when $M \geq n_A \cdot n_B$, once a match is severed, it is never reactivated again. Therefore, matches are gradually broken over time both under profit and under welfare maximization. In this case, all relationships are severed too early (weakly) under profit maximization.

This last property, however, does not extend to environments with binding capacity constraints. Suppose, for example, that $M = 1$ and that processes are endogenous. For any pair of agents $(i, j) \in N_A \times N_B$ and any history of past interactions between the pair, the index assigned to the pair in a profit-maximizing auction is always smaller than the index in the corresponding welfare-maximizing auction. Formally, this can be seen by noticing that the weights $\hat{\beta}$ used to compute the indexes under profit maximization are strictly smaller than the corresponding weights in a welfare-maximizing auction. However, the ranking of the indexes across pairs of agents under profit maximization need not coincide with the ranking under welfare maximization. As a result, certain interactions may last longer under profit maximization than under welfare maximization. What remains true, though, is that, if at a given point in time matching shuts down under welfare maximization, then, under profit maximization, either matching shut down already in previous periods, or it shuts down in the present period.

Similarly, when processes are exogenous, for any M , the aggregate level of interactions at any period under profit maximization is always (weakly) lower than under welfare maximization, although some individual interactions may last longer under profit maximization.

When, instead, processes are endogenous and the capacity constraint takes intermediate levels (that is, $1 < M < n_A \cdot n_B$), welfare maximization does not necessarily induce more aggregate activity in each period than profit maximization; nor is it necessarily the case that, intertemporally, more matches are implemented under welfare maximization than under profit maximization. This latter property is guaranteed only if matching terminates in finite time under profit maximization. We refer the reader to the Supplementary Material for an example in which the number of matches is always weakly larger under profit maximization than under welfare maximization (and strictly larger in certain periods).

Interestingly, when certain agents dislike certain interactions (formally, when horizontal types may take on negative values for certain pairs), the conclusions in Theorem 4 need not hold, and profit maximization may distort matching in the opposite direction. That is, profit-maximizing auctions may induce an inefficiently *high* volume of matches within each period, and interactions may last longer under profit maximization than under welfare maximization, *irrespective of the capacity constraint*. Formally, when certain interactions may generate negative match values, a pair’s myopic score under profit maximization may be greater than its counterpart under welfare maximization. As a result, the conclusions in the above theorem can be overturned. Intuitively, the reason why a profit-maximizing platform may induce an inefficiently high volume of interactions is that this may permit it to economize on the informational rents that it must leave to the agents. By locking those agents selecting a low membership status (equivalently, claiming to have a low vertical type) into unpleasant interactions, the platform makes it costly for those agents with a high vertical type to pretend to have a low type. In turn, this permits the platform to extract more surplus from those high-type agents. The Supplementary Material contains an example illustrating this point.

Finally, note that the familiar result of “no distortion at the top” from standard screening problems does not apply to a matching environment. A profit-maximizing platform may distort the matches of all agents, including

those “at the top” of the distribution, for whom the vertical type is the highest. The reason is that, contrary to standard screening problems in which the cost of procuring inputs is exogenous, in a matching market, the cost of “procuring” agents-inputs from the opposite side of the market is endogenous and is higher than under welfare maximization, due to the informational rents that the platform must provide to such agents-inputs to induce them to reveal their private information.

The above results have implications for the design of government intervention in matching markets, a topic that is receiving a lot of attention in the policy debate. The results in Theorem 3 indicate that, in many markets of interest, the government could impose that the scoring rules be tilted by adopting the welfare-maximizing weights β^W in lieu of the profit-maximizing ones $\hat{\beta}$. However, because the welfare-maximizing auctions are not guaranteed to yield positive profits (even under the profit-maximizing payments of part (iii) in Theorem 3), the government may need to subsidize the platform.

The results in Theorem 4, in turn, suggest that simple subsidies aimed at inducing platforms to increase the volume of matches they accommodate in each period can be welfare-increasing in markets with non-binding capacity constraints, but might be counterproductive when such constraints bind. In this latter case, regulators may need to target the subsidies to specific matches, which may be feasible in markets with a small number of agents, but seems difficult in large anonymous markets.

Lastly, the above results indicate that, when certain agents dislike certain interactions, the government may want to tax certain interactions while subsidizing others. Again, whether this is at all possible is likely to depend on how anonymous the interactions in such markets are.

8 Conclusions

This paper introduces and then studies a class of auctions that platforms can use to match agents from different sides of the market in environments in which agents’ preferences for potential partners are their own private information and evolve over time, either exogenously, or as a function of previous interactions.

The proposed auctions are fairly simple and can be used in a variety of markets including those for scientific outsourcing, peer-to-peer lending, ad exchanges, and organized events. Upon joining the platform, agents select a membership status which determines the weight assigned to their bids in the subsequent auctions. They then bid repeatedly for potential partners from the opposite side of the market. The platform then computes bilateral scores for each match and implements the matches with the highest nonnegative score, up to capacity.

In case the processes governing the evolution of the match values are exogenous, the scores are myopic and reflect the flow values the agents assign to the matches, net of the platform’s costs of providing all the auxiliary services necessary to implement the matches, and net of handicaps controlling for the agents’ information rents. When, instead, the match values depend on past interactions, the scores are forward-looking and control for the value of generating information that can be used in subsequent periods. Under appropriate conditions, such forward-looking scores take the form of indexes similar to those in the operation research literature (e.g., Gittins, 1979), but adjusted to account for the costs of information rents.

The framework is flexible enough to admit as special cases such environments in which learning occurs immediately upon matching, as well as such environments in which agents enjoy interacting at most finitely

many times with each partner. The results are then used to shed light on the inefficiencies associated with the private provision of matching services, which in turn can be used to evaluate the merits of certain government interventions.

To the best of our knowledge, the auctions we propose are not in use in any specific platform market. However, we believe they capture the key trade-offs that many platforms face in the design of their dynamic matching protocols (see, for example, the discussion in Chapter 6 in Moghaddam and Shimon (2016)).

Many extensions appear interesting. Certain results can be adapted to accommodate for the possibility that match values depend on the entire history of past interactions (e.g., agents may care about their partners' previous partners). Extending the analysis to allow for more general forms of correlation in the agents' preferences is challenging but worth exploring. It also appears interesting to study how the results specialize in markets in which formal payments cannot be collected from one side of the market (as with certain types of sponsored search).

It would also be interesting to consider a broader class of preferences and processes governing the evolution of the match values. The ones considered in the present paper favor tractability. As noticed above, no general solutions are known in the experimentation literature for multi-armed bandit problems in which multiple arms are activated simultaneously. However, there are environments in which asymptotic results can be established for large markets. For example, Bergemann and Valimaki (2001) show that an index policy is optimal for stationary multi-armed bandit problems, in which there are countable infinitely many ex-ante identical arms (and approximately optimal in the limit as the number of arms goes to infinity). For the more general case of *restless bandits* (which involves arms that evolve and yield rewards even when they are not activated), under the infinite-horizon average reward criterion, Weber and Weiss (1990) provide conditions guaranteeing that an index policy that selects in each period the arms with the highest Whittle index is asymptotically optimal (see also Whittle (1988)).³⁸ Such asymptotic results could also be useful in designing scoring rules for large matching markets. It would also be interesting to examine the design of dynamic matching protocols with simple scores in environments in which the processes governing the evolution of the match values are not known and where, as a result, platforms must use simple heuristics to trade-off the value of exploitation with the value of experimentation.

In the same vein, it would be interesting to consider a variant in which the platform must incur a cost only when it matches a pair of agents who were not matched in the preceding period. Such costs introduce an interesting form of inertia and give rise to a trade-off between improving the quality of the existing matches and economizing on costs that is absent in our analysis.

Another fruitful direction for future research consists in comparing the matching dynamics in centralized markets such as the ones investigated in the present paper to their counterparts in decentralized markets. To the best of our knowledge, there is no tractable model of decentralized matching where agents perform on-the-job search and rematch over time, in environments with exogenously or endogenously evolving match values. Developing such a model is challenging but an important next step for this literature.

We conclude by noting that, while matching dynamics in the present paper originate in changes in agents' preferences for potential partners, another line of recent research explores matching dynamics driven by the

³⁸In the special case in which passive arms are static and yield no reward, this index reduces to the Gittins index.

(stochastic) arrival or departure of agents to and from the market (see, for example, Anderson et al. 2015, Baccara et al. 2015, and Akbarpour et al. 2016). Incorporating stochastic arrivals and departures into the scores of the matching auctions introduced in the present paper is another interesting direction for future research.

Appendix

Proof of Theorem 1. The proof is in three steps. Step 1 shows that, when all agents follow truthful strategies, in each of the respective environments of Theorem 1, the matches under the rules $\chi^{m;\beta}$ and $\chi^{I;\beta}$ maximize the continuation weighted surplus, as defined in (6), starting from any period- t history, any $t \geq 1$. Step 2 shows that when, in addition, the period- t payments are as in (8), any $t \geq 1$, then participating in the auctions and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period- t history, any $t \geq 1$. Finally, Step 3 completes the proof by showing that the matching rules $\chi^{m;\beta}$ and $\chi^{I;\beta}$ satisfy a certain dynamic monotonicity condition (defined below), which guarantees that when, in addition, the period-0 membership fees are as in (9), then participating and following truthful strategies is a periodic ex-post equilibrium starting from period zero.

Step 1. We first establish that, in each of the respective environments of Theorem 1, the matching rules $\chi^{m;\beta}$ and $\chi^{I;\beta}$ maximize continuation weighted surplus (6). That is, irrespective of the particular history that led to the selection of the period-0 membership statuses θ_0 and of the past matches x^{t-1} , in the continuation game that starts with period t , $t \geq 1$, when the true vertical type profile is θ_t , the true profile of horizontal types is ε_t (with ε_t obtained from θ_t and b_t using (4)), and when agents follow truthful strategies from period t onwards, the matches under $\chi^{m;\beta}$ and $\chi^{I;\beta}$ maximize W_t , over the entire set \mathcal{X} of all feasible matching rules, with message spaces defined as in the matching auctions (i.e., $\mathcal{M}_{l0}^k = \Theta_l^k$, and $\mathcal{M}_{lt}^k = \Theta_l^k \times \mathbb{R}^{N-k}$, $l \in N_k$, $k = A, B$, $t \geq 1$).

Lemma 1 below considers the case of exogenous processes and the case of endogenous processes with extreme capacity constraints ($M = 1$ and $M \geq n_A \cdot n_B$). Proposition 1 below considers the case of endogenous processes with intermediate capacity constraints, under the separability assumption.

Lemma 1 (i) *Suppose the processes are exogenous. Then $\chi^{m;\beta}$ maximizes continuation weighted surplus at all histories.* (ii) *Suppose processes are endogenous and either $M \geq n_A \cdot n_B$, or $M = 1$. Then $\chi^{I;\beta}$ maximizes continuation weighted surplus at all histories.*

Proof of Lemma 1. Part (i) follows directly from the fact that $\chi^{m;\beta}$ maximizes (6) history by history. For part (ii), note that the problem of maximizing (6) can be viewed as a multiarmed bandit problem, with each arm corresponding to a potential match, and with the flow period- t reward of activating each arm (i, j) given by the myopic score $S_{ijt}^{m;\beta}$. When $M = 1$, that $\chi^{I;\beta}$ maximizes (6) at all histories is then immediate, as the score $S_{ijt}^{I;\beta}$ corresponds to the arm's Gittins Index (see, for example, Whittle (1982)). Similarly, when $M \geq n_A \cdot n_B$, since the capacity constraint never binds, the platform's problem can be viewed as a collection of $n_A \cdot n_B$ separate two-armed bandit problems, one for each potential pair of agents, with the reward from

matching the pair (i, j) given by $S_{ijt}^{m;\beta}$ and the reward from activating the “safe arm” identically equal to zero. In both cases, $\chi^{I;\beta}$ maximizes continuation weighted surplus. ■

Next, consider the case of endogenous processes and arbitrary capacity constraints, under the separability assumption.

Proposition 1 *Suppose processes are endogenous and the environment is separable under the rule $S^{I;\beta}$. For any $M \in \mathbb{N}$, the matching rule $\chi^{I;\beta}$ maximizes continuation weighted surplus (6) starting from any period- t history (θ^t, b^t, x^{t-1}) , any $t \geq 1$.*

Proof of Proposition 1. The proof follows from the four lemmas below. Consider a fictitious environment in which the rewards are given by the auxiliary processes defined in (11). That is, for any $t \geq 1$, $(i, j) \in N_A \times N_B$, (θ^t, b^t, x^{t-1}) , suppose the period- t reward from the match (i, j) is given by $\underline{S}_{ijt}^\beta$. Denote by $\underline{\chi}^{m;\beta}$ the myopic matching rule that, in each period, matches the pairs with the highest nonnegative auxiliary rewards $\underline{S}_{ijt}^\beta$, subject to the platform’s capacity constraint (ties broken arbitrarily). Formally, let

$$\underline{Q}_t^\beta(\theta^t, b^t, x^{t-1}) \equiv \left\{ \begin{array}{l} (i, j) \in N_A \times N_B \text{ s.t. (i) } \underline{S}_{ijt}^\beta(\theta^t, b^t, x^{t-1}) \geq 0 \text{ and} \\ \text{(ii) } \# \left\{ (l, m) \in N_A \times N_B : \underline{S}_{lmt}^\beta(\theta^t, b^t, x^{t-1}) > \underline{S}_{ijt}^\beta(\theta^t, b^t, x^{t-1}) \right\} < M \end{array} \right\}.$$

Then (a) $\underline{\chi}_{ijt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 1$ only if $(i, j) \in \underline{Q}_t^\beta(\theta^t, b^t, x^{t-1})$. Furthermore, (b) if $\underline{S}_{ijt}^\beta(\theta^t, b^t, x^{t-1}) > 0$ and $\# \left\{ (l, m) \in N_A \times N_B : \underline{S}_{lmt}^\beta(\theta^t, b^t, x^{t-1}) \geq \underline{S}_{ijt}^\beta(\theta^t, b^t, x^{t-1}) \right\} \leq M$, then $\underline{\chi}_{ijt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 1$. Finally, (c), if $(i, j) \in \underline{Q}_t^\beta(\theta^t, b^t, x^{t-1})$, $\underline{S}_{ijt}^\beta(\theta^t, b^t, x^{t-1}) > 0$, and also $\underline{\chi}_{ijt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 0$, then it must be the case that $\# \left\{ (l, m) \in N_A \times N_B : \underline{\chi}_{lmt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 1 \right\} = M$.

Lemma 2 *Suppose processes are endogenous and the environment is separable under the rule $S^{I;\beta}$. Then in the corresponding fictitious environment in which the flow rewards are given by \underline{S}^β , the rule $\underline{\chi}^{m;\beta}$ maximizes the expected discounted sum of the auxiliary rewards starting from any period- t history, any $t \geq 1$.³⁹*

Proof of Lemma 2. Suppose, towards a contradiction, that the claim is not true. This means that there exists a period $t \geq 1$, a history (θ^t, b^t, x^{t-1}) , and a feasible rule $\underline{\chi} \neq \underline{\chi}^{m;\beta}$ such that the expected discounted sum of auxiliary rewards from period t onwards is higher under $\underline{\chi}$ than under $\underline{\chi}^{m;\beta}$. For this to be the case, there must exist a period $s \geq t$, and a set of histories (θ^s, b^s, x^{s-1}) of strictly positive probability under $\lambda[\underline{\chi}]|\theta^t, b^t, x^{t-1}$, for which the matches under the two rules differ, meaning that one of the following two properties (or both) must hold, for some $(i, j) \in N_A \times N_B$: (a) either $\underline{\chi}_{ijs}(\theta^s, b^s, x^{s-1}) = 1$ and $(i, j) \notin \underline{Q}_s^\beta(\theta^s, b^s, x^{s-1})$, in which case $\underline{\chi}_{ijs}^m(\theta^s, b^s, x^{s-1}) = 0$; (b) or $\underline{\chi}_{ijs}(\theta^s, b^s, x^{s-1}) = 0$, $\underline{\chi}_{ijs}^{m;\beta}(\theta^s, b^s, x^{s-1}) = 1$, $\underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1}) > 0$, and

$$\# \left\{ (l, m) \in N_A \times N_B : \underline{S}_{lms}^\beta(\theta^s, b^s, x^{s-1}) \geq \underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1}), \underline{\chi}_{lms}(\theta^s, b^s, x^{s-1}) = 1 \right\} < M.$$

Note that any other case in which the two rules $\underline{\chi}^m$ and $\underline{\chi}$ implement different allocations for the pair (i, j) and neither (a) nor (b) are satisfied is inconsequential for the difference in the expected continuation weighted

³⁹The formula for the expected discounted sum of the auxiliary rewards is the same as the one in (6), but with \underline{S}_{ijs} replacing S_{ijs} , and each χ_{ijs} defined over (θ^s, b^s, x^{s-1}) as opposed to $(\theta_0, \theta_s, b_s, x^{s-1})$, all $(i, j) \in N_A \times N_B$, $s \geq t$. As in the case of the formula in (6), the expectation over future values and types is under truthful strategies.

surplus under the two rules (in fact, the difference in the allocations in any such other case simply reflects the way ties are broken).

First, consider case (a). Because $(i, j) \notin Q_s^\beta(\theta^s, b^s, x^{s-1})$, either (a1) $\underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1}) < 0$ or (a2)

$$\# \left\{ (l, m) \in N_A \times N_B : \underline{S}_{lms}^\beta(\theta^s, b^s, x^{s-1}) > \underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1}) \right\} \geq M.$$

In the first case (case a1), given that the auxiliary rewards are non-increasing, it is immediate that $\underline{\chi}^{m;\beta}$, by leaving the pair (i, j) unmatched in period s , improves upon $\underline{\chi}$. Thus consider the second possibility, case (a2). There must exist another pair $(i', j') \in N_A \times N_B$ such that $\underline{S}_{i'j's}^\beta(\theta^s, b^s, x^{s-1}) > \underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1}) \geq 0$, $\underline{\chi}_{i'j's}(\theta^s, b^s, x^{s-1}) = 0$ and $\underline{\chi}_{i'j's}^m(\theta^s, b^s, x^{s-1}) = 1$. Because the environment is separable under the rule $S^{I;\beta}$, this implies that $\underline{S}_{i'j's}^\beta(\theta^s, b^s, x^{s-1}) \geq \underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1})/(1 - \delta)$. By definition, the auxiliary processes are non-increasing. This means that the expected discounted sum of the stream of auxiliary rewards that can be obtained by matching (i, j) in period s (and possibly in some of the subsequent periods) is no greater than $\underline{S}_{ijs}^\beta(\theta^s, b^s, x^{s-1})/(1 - \delta)$, which is smaller than the flow reward $\underline{S}_{i'j's}^\beta(\theta^s, b^s, x^{s-1})$ obtained by matching (i', j') in period s . Hence, by favoring (i', j') over (i, j) in period t , $\underline{\chi}^{m;\beta}$ again improves upon the rule $\underline{\chi}$.

Next, consider case (b). The myopic rule $\underline{\chi}^{m;\beta}$ improves upon $\underline{\chi}$ by either adding the match (i, j) to the set of matches implemented under $\underline{\chi}$ (in case $\underline{\chi}$ matches fewer than M pairs) or by matching the pair (i, j) instead of another pair (i', j') that is matched under the rule $\underline{\chi}$ but not under the rule $\underline{\chi}^{m;\beta}$ and for which the flow auxiliary reward is strictly smaller than (i, j) 's (in which case the separability assumption again guarantees that such a change improves upon $\underline{\chi}$).

In either cases (a) and (b) above, $\underline{\chi}^{m;\beta}$ thus improves upon $\underline{\chi}$. Applying the arguments above to all histories and all pairs of agents yields a contradiction to the fact that $\underline{\chi}$ strictly dominates $\underline{\chi}^{m;\beta}$. ■

The next two lemmas fix the matching rule and relate the expected discounted sum of the rewards in the primitive environment (where the flow rewards are given by $S^{m;\beta}$) to the expected discounted sum of the rewards in the fictitious environment (where the rewards are the auxiliary ones, \underline{S}^β). The proof of the next two lemmas is omitted, for it follows from arguments similar to those in Weber (1992) (see also Mandelbaum, 1986, Ishikida and Varaiya 1994, and Pandelis and Teneketzis, 1999), adapted to the matching environment under examination.

Lemma 3 *Suppose the agents follow truthful strategies. The (period-0) expected discounted sum of the auxiliary rewards \underline{S}^β under any rule χ is weakly higher than the (period-0) expected discounted sum of the primitive rewards $S^{m;\beta}$ under the same rule χ .*

To understand the result, note first that the auxiliary rewards are themselves Gittins indexes, and are therefore defined selecting the stopping times that maximize the expected average discounted payoff per unit of expected discounted time. Also note that the optimal stopping times coincide with the first time at which the index drops weakly below its initial value. Therefore, under any arbitrary matching rule where the stopping times are possibly different from the optimal ones, the expected discounted sum of the true rewards $S^{m;\beta}$ can never exceed the expected discounted sum of the auxiliary rewards \underline{S}^β .

In the special case where the matching rule is the myopic rule for the auxiliary rewards, $\underline{\chi}^{m;\beta}$, the stopping

times implemented under such rule coincide with those under the Gittins policy for the primitive rewards. As a result, the expected discounted sum of the primitive rewards is the same as that for the auxiliary rewards.

Lemma 4 *Suppose the agents follow truthful strategies. The (period-0) expected discounted sum of the auxiliary rewards \underline{S}^β under the myopic rule $\underline{\chi}^{m;\beta}$ for the auxiliary processes is the same as the (period-0) expected discounted sum of the primitive rewards $S^{m;\beta}$ under the same rule $\underline{\chi}^{m;\beta}$.*

Finally, the following property is a direct consequence of the definition of the index rule $\chi^{I;\beta}$ for the primitive environment.

Lemma 5 *Suppose the agents follow truthful strategies. The dynamics of the matches under $\underline{\chi}^{m;\beta}$ are the same as under $\chi^{I;\beta}$.*

The proof of Proposition 1 then follows from combining the results in the above four lemmas. ■

Step 2. Fix the weights β and denote by $\tilde{\chi}$ a matching rule that maximizes the continuation weighted surplus at all histories, and by $\tilde{\chi}_{-l}^k$ a matching rule that does so in the absence of agent l from side $k \in \{A, B\}$ (equivalently, that maximizes continuation weighted surplus when the myopic score of any match that involves agent l from side k is identically equal to zero, in which case $\tilde{\chi}_{-l}^k$ can be assumed to never implement any match involving agent l). The existence of such rules, in the corresponding environments of Theorem 1, has been shown in Step 1 above.⁴⁰ Denote by $\tilde{\psi}_{t \geq 1} \equiv (\tilde{\psi}_s)_{s \geq 1}$ the collection of payment functions, given by (8), starting from period one, defined with respect to the matching rule $\tilde{\chi}$. Henceforth, the weighted surpluses W_t and $W_t^{-l,k}$, as well as the resulting marginal and flow contributions to weighted surplus, as defined in the main text, unless otherwise specified, are with respect to the matching rules $\tilde{\chi}$ and $\tilde{\chi}_{-l}^k$, respectively. Note that, because the weights β are held fixed, to ease the notation we drop them from all functions below, when there is no risk of confusion.

Lemma 6 *Consider an auction in which the matching rule is $\tilde{\chi}$ and the payments from period 1 onwards are given by $\tilde{\psi}_{t \geq 1}$. In such an auction, participating and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period- t history, $t \geq 1$.*

Proof of Lemma 6. We show that, in the continuation game that starts in period $t \geq 1$, irrespective of the history of past play, of the true vertical type profile θ , and of the history of past and current horizontal types $\varepsilon^t \equiv (\varepsilon_s)_{s=1}^t$, any agent $l \in N_k$, $k = A, B$, who expects all other agents to participate and follow truthful strategies from period t (included) onwards, finds it optimal to do the same.

Consider agent l from side A (the problem for any agent from side B is similar). Suppose that the true profile of vertical types is θ , the true profile of period- t match values is u_t , the profile of period-0 membership choices is θ_0 , and the history of past matches is x^{t-1} . Denote by

$$\mathbb{E}^{\tilde{\chi}[\tilde{\chi}]|\theta, v_t, x^{t-1}; (\hat{\theta}_{it}^A, \hat{b}_{it}^A)} [R_{it+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t)]$$

⁴⁰These rules correspond to the myopic and index rules for exogenous and endogenous processes, respectively.

the expected contribution of agent l to continuation weighted surplus from period $t+1$ onwards, when, in period t , the agent selects the period- t membership status $\hat{\theta}_{lt}^A$, submits the period- t bids \hat{b}_{lt}^A , follows a truthful strategy from period $t+1$ onwards, and expects all other agents to follow truthful strategies at all periods $s \geq t$.⁴¹ Note that, when the agent follows the truthful strategy also in period t (i.e., when $\hat{\theta}_{lt}^A = \theta_l^A$ and $\hat{b}_{lt}^A = v_{lt}^A$), then

$$\tilde{\lambda}[\tilde{\chi}]|\theta, v_t, x^{t-1}; (\theta_l^A, v_{lt}^A) = \lambda[\tilde{\chi}]|\theta, v_t, x^{t-1},$$

where the process $\lambda[\tilde{\chi}]|\theta, v_t, x^{t-1}$ is as defined in the main text.

Since the agent can adjust his membership status in any of the subsequent periods, any deviation from the truthful strategy in period t can be corrected in period $t+1$. This means that, to prove the result, it suffices to show that the agent prefers to follow the truthful strategy from period t onwards than deviating in period t and then reverting to the truthful strategy from period $t+1$ onwards.

Under the proposed auction rules, when the agent follows the truthful strategy from period $t+1$ onwards, his continuation payoff from period $t+1$ onwards is given by

$$\frac{1}{\beta_l^A(\theta_{l0}^A)} R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t)$$

Therefore, it is enough to show that, for any period- t selection $(\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)$,

$$\begin{aligned} & \sum_{j \in N_B} v_{ljt}^A \tilde{\chi}_{ljt}(\theta_0, (\theta_l^A, \theta_{-l}^A), (v_{lt}^A, v_{-lt}^A), x^{t-1}) - \tilde{\psi}_{lt}^A(\theta_0, (\theta_l^A, \theta_{-l}^A), (v_{lt}^A, v_{-lt}^A), x^{t-1}) \\ & + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}} [R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t)] \\ & \geq \sum_{j \in N_B} v_{ljt}^A \tilde{\chi}_{ljt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) - \tilde{\psi}_{lt}^A(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\ & + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}]|\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} [R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t)]. \end{aligned} \quad (16)$$

The left hand side of the above inequality can be rewritten in terms of the functions W_t and $W_t^{-l,A}$ as follows:

$$\frac{1}{\beta_l^A(\theta_{l0}^A)} \left[W_t(\theta_0, (\theta_l^A, \theta_{-l}^A), (v_{lt}^A, v_{-lt}^A), x^{t-1}) - W_t^{-l,A}(\theta_0, (\theta_l^A, \theta_{-l}^A), (v_{lt}^A, v_{-lt}^A), x^{t-1}) \right]. \quad (17)$$

That is, agent l 's expected continuation payoff when he follows the truthful strategy from period t onward is equal to his expected contribution to the maximal continuation weighted surplus, scaled by the weight $\beta_l^A(\theta_{l0}^A)$. It then suffices to show that (17) is weakly greater than the right hand side of (16).

⁴¹The stochastic process $\tilde{\lambda}[\tilde{\chi}]|\theta, \varepsilon_t, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)$ is here over future matches, bids and membership choices, under the rule $\tilde{\chi}$, when agent l 's period- t choices are $(\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)$, the profile of vertical types is θ , the true profile of period- t horizontal types is ε_t , the history of past matches is x^{t-1} , the agent plans to follow a truthful strategy from $t+1$ onwards and all other agents follow truthful strategies from period t onwards.

Next, note that the flow contribution r_{lt}^k can be rewritten as

$$\begin{aligned}
r_{lt}^k(\theta_0, \theta_t, b_t, x^{t-1}) &= \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}(\theta_0, \theta_t, b_t, x^{t-1}) \cdot \tilde{\chi}_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) \\
&\quad - \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}(\theta_0, \theta_t, b_t, x^{t-1}) \tilde{\chi}_{ijt}^{-l,k}(\theta_0, \theta_t, b_t, x^{t-1}) \\
&\quad + \delta \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_t, b_t, x^{t-1}} \left[W_{t+1}^{-l,k}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
&\quad - \delta \mathbb{E}^{\lambda[\tilde{\chi}^{-l,k}]|\theta_t, b_t, x^{t-1}} \left[W_{t+1}^{-l,k}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right].
\end{aligned} \tag{18}$$

Using (18), we can rewrite the period- t payment in the right-hand side of (16) as follows:

$$\begin{aligned}
&\tilde{\psi}_{lt}^A(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&= \sum_{j \in N_B} \hat{b}_{ljt}^A \tilde{\chi}_{ljt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) - \frac{1}{\beta_l^A(\theta_{l0}^A)} r_{lt}^A(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&= \sum_{j \in N_B} \hat{b}_{ljt}^A \tilde{\chi}_{ljt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&\quad - \frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \cdot \tilde{\chi}_{ijt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&\quad - \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_{lt}^A, \theta_{-l}^A, \hat{b}_{lt}^A, v_{-lt}^A, x^{t-1}} \left[W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
&\quad + \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&= -\frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N_A \setminus \{l\}} \sum_{j \in N_B} S_{ijt}^{m;\beta}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \cdot \tilde{\chi}_{ijt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&\quad - \frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{j \in N_B} (\beta_j^B(\theta_{j0}^B) v_{ljt}^B - c_{ljt}(x^{t-1})) \cdot \tilde{\chi}_{ljt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\
&\quad - \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_{lt}^A, \theta_{-l}^A, \hat{b}_{lt}^A, v_{-lt}^A, x^{t-1}} \left[W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
&\quad + \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}).
\end{aligned}$$

Furthermore, note that

$$\begin{aligned}
&\mathbb{E}^{\lambda[\tilde{\chi}]|\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
&= \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) - W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
&= \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
&\quad - \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_{lt}^A, \theta_{-l}^A, \hat{b}_{lt}^A, v_{-lt}^A, x^{t-1}} \left[W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right],
\end{aligned}$$

where the last equality uses the fact that, given x^t , $W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)$ is invariant in agent l 's period- t bids and that the period- t decisions x_t are invariant in the agent's true types. Therefore, the right hand side

of the inequality (16) is equal to

$$\begin{aligned} & \frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}(\theta_0, \theta, v_t, x^{t-1}) \cdot \tilde{\chi}_{ijt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\ & + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\chi}}[\tilde{\chi}]_{\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} [W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)] \\ & - \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}). \end{aligned}$$

Since, again, $W_t^{-l,A}$ is independent of agent l 's period- t bids, to establish that the inequality in (16) holds, it suffices to show that

$$\begin{aligned} W_t(\theta_0, \theta, v_t, x^{t-1}) & \geq \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}(\theta_0, \theta, v_t, x^{t-1}) \cdot \tilde{\chi}_{ijt}(\theta_0, (\hat{\theta}_{lt}^A, \theta_{-l}^A), (\hat{b}_{lt}^A, v_{-lt}^A), x^{t-1}) \\ & + \mathbb{E}^{\tilde{\chi}}[\tilde{\chi}]_{\theta_l^A, \theta_{-l}^A, v_{lt}^A, v_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} [W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)]. \end{aligned} \quad (19)$$

The inequality in (19) follows from the definition of the matching rule $\tilde{\chi}$.

That it is (periodic ex-post) optimal for each agent to participate at all periods $t \geq 1$, and after all histories, follows from the fact that each agent's continuation payoff under truthful strategies coincides with his expected contribution to continuation weighted surplus, which is always nonnegative, scaled by a strictly positive weight.

The arguments above therefore establish that, when the matching rule is $\tilde{\chi}$ and the payments from period 1 onwards are $\tilde{\psi}_{\geq 1}$, participating and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period- t history, $t \geq 1$. ■

Step 3. We now show that, when the period-0 membership fees are as in (9), participating in period zero and then following a truthful strategy at all periods (including period zero) is a periodic ex-post equilibrium.

Let

$$\tilde{D}_l^k(\theta_{l0}^k, \theta; \chi) \equiv \begin{cases} \mathbb{E}^{\lambda[\chi]|\theta} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N_{-k}} \varepsilon_{lht}^k \chi_{lht}(\theta_0, \theta_t, b_t, x^{t-1}) \right] & \text{if } k = A \\ \mathbb{E}^{\lambda[\chi]|\theta} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N_{-k}} \varepsilon_{hlt}^k \chi_{hlt}(\theta_0, \theta_t, b_t, x^{t-1}) \right] & \text{if } k = B \end{cases},$$

denote the match quality that agent $l \in N_k$ from side $k = A, B$ expects under the rule χ when the true profile of vertical types is $\theta \in \Theta$, the agent selects the membership status θ_{l0}^k in period zero and then conforms to the truthful strategy from period $t = 1$ onwards, and all agents other than l (from side k) follow truthful strategies at each period. Note that $\lambda[\chi]|\theta$ denotes the stochastic process over matches, bids, and membership choices, when the true vertical types are θ , and all agents follow truthful strategies from period $t = 1$ onwards. Also note that $\tilde{D}_l^k(\theta_l^k, (\theta_{-l}^k, \theta_l^k); \chi) = D_l^k(\theta; \chi)$, with the function $D_l^k(\theta; \chi)$ as defined in (10) — hereafter we highlight the dependence of the function $D_l^k(\theta; \chi)$ on the matching rule χ to avoid possible confusion.

For any agent $i \in N_A$ (the arguments for the side- B agents are analogous), let

$$\hat{U}_i^A(\theta) \equiv \mathbb{E}^{\lambda[\chi]|\theta} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} \theta_i^A \varepsilon_{ijt}^A \chi_{ijt}(\theta, \theta_t, b_t, x^{t-1}) \right] - \mathbb{E}^{\lambda[\chi]|\theta} \left[\sum_{t=0}^{\infty} \delta^t \psi_{it}^A(\theta, \theta_t, b_t, x^{t-1}; \beta) \right]$$

denote the payoff that the agent expects in the matching auctions defined by the rules (χ, ψ) when the true

vertical type profile is θ and all agents follow truthful strategies at all periods.⁴²

Let $\tilde{\chi}$ be a matching rule that maximizes continuation weighted surplus and $\tilde{\psi} = (\tilde{\psi}_0, \tilde{\psi}_{\geq 1})$ the associated payment rule, as defined in the main text. Note that, for each agent $l \in N_k$, each profile θ of true vertical types,

$$\mathbb{E}^{\lambda[\tilde{\chi}]|\theta} \left[\tilde{\psi}_{l0}^k(\theta; \beta) + \sum_{t=1}^{\infty} \delta^t \tilde{\psi}_{lt}^k(\theta, \theta_t, b_t, x^{t-1}; \beta) \right] = \theta_l^k D_l^k(\theta; \tilde{\chi}) - \int_{\underline{\theta}_l^k}^{\theta_l^k} D_l^k(\theta_{-l}^k, y; \tilde{\chi}) dy - L_l^k,$$

which guarantees that, when all agents follow truthful strategies in each period, including period zero, the period-zero expected payoffs, $\hat{U}_l^k(\theta)$, are given by

$$\hat{U}_l^k(\theta) = \int_{\underline{\theta}_l^k}^{\theta_l^k} D_l^k(\theta_{-l}^k, y; \tilde{\chi}) dy + L_l^k. \quad (20)$$

The next lemma shows that any matching rule $\tilde{\chi}$ that maximizes continuation weighted surplus (6) at all histories satisfies a certain monotonicity condition which plays a central role in establishing the optimality of the truthful strategies at period zero.

Lemma 7 *Suppose $\tilde{\chi}$ maximizes continuation weighted surplus (6) at all histories. For all $l \in N_k$, $k = A, B$, the following monotonicities hold:*

- (i) $D_l^k((\theta_{-l}^k, \theta_l^k); \tilde{\chi})$ is non-decreasing in θ_l^k , all $\theta_{-l}^k \in \Theta_{-l}^k$;
- (ii) $\tilde{D}_l^k(\theta_{l0}^k, \theta; \tilde{\chi})$ is non-decreasing in θ_{l0}^k , all $\theta \in \Theta$.

Proof of Lemma 7. Consider an arbitrary agent $i \in N_A$ from side A (the arguments for the side- B agents are analogous) and fix the profile of types θ_{-i}^A for the other agents.

Part (i). Take any pair of types $\theta_i^A, \hat{\theta}_i^A \in \Theta_i^A$, with $\theta_i^A < \hat{\theta}_i^A$. That $\tilde{\chi}$ maximizes continuation weighted surplus implies that⁴³

$$\begin{aligned} & \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N_A \setminus \{i\}} \sum_{j \in N_B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}(x^{t-1})) \tilde{\chi}_{rjt}((\theta_{-i}^A, \hat{\theta}_i^A), (\theta_{-i}^A, \hat{\theta}_i^A), b_t, x^{t-1}) \right] \\ & + \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} (\beta_i^A(\hat{\theta}_i^A) b_{ijt}^A + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}(x^{t-1})) \tilde{\chi}_{ijt}((\theta_{-i}^A, \hat{\theta}_i^A), (\theta_{-i}^A, \hat{\theta}_i^A), b_t, x^{t-1}) \right] \\ & \geq \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N_A \setminus \{i\}} \sum_{j \in N_B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_r^B(\theta_r^B) b_{rjt}^B - c_{rjt}(x^{t-1})) \tilde{\chi}_{rjt}((\theta_{-i}^A, \theta_i^A), (\theta_{-i}^A, \theta_i^A), b_t, x^{t-1}) \right] \\ & + \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} \left(\beta_i^A(\hat{\theta}_i^A) b_{ijt}^A \frac{\hat{\theta}_i^A}{\theta_i^A} + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}(x^{t-1}) \right) \tilde{\chi}_{ijt}((\theta_{-i}^A, \theta_i^A), (\theta_{-i}^A, \theta_i^A), b_t, x^{t-1}) \right]. \end{aligned}$$

The left-hand side of the previous inequality is the expected weighted surplus when all agents follow truthful strategies from period $t = 0$ onward and the true profile of vertical types is $(\hat{\theta}_i^A, \theta_{-i}^A)$. The right-hand side is the expected weighted surplus when, under the same profile of true vertical types $(\hat{\theta}_i^A, \theta_{-i}^A)$, all agents other than

⁴²The dependence of the payoff on the mechanism (χ, ψ) is omitted for convenience.

⁴³Recall that $\lambda[\tilde{\chi}]|\theta_i^A, \theta_{-i}^A$ denotes the process over matches, bids, and membership choices from period $t = 1$ onwards, when all agents follow truthful strategies from period zero onwards, and the true vertical types are $(\theta_i^A, \theta_{-i}^A)$.

agent i from side A follow truthful strategies in all periods whereas agent i follows the strategy of type θ_i^A in all periods (that is, at each period he selects the membership status $\theta_{it}^A = \theta_i^A$ and then submits bids equal to $b_{ijt}^A = \theta_i^A \varepsilon_{ijt}^A$, where ε_{ijt}^A are the true horizontal types).

Similarly, inverting the role of $\hat{\theta}_i^A$ and θ_i^A , we have that

$$\begin{aligned} & \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N_A \setminus \{i\}} \sum_{j \in N_B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}(x^{t-1})) \tilde{\chi}_{rjt}((\theta_{-i}^A, \theta_i^A), (\theta_{-i}^A, \theta_i^A), b_t, x^{t-1}) \right] \\ & + \mathbb{E}^{\lambda[\tilde{\chi}]|\theta_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} (\beta_i^A(\theta_i^A) b_{ijt}^A + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}(x^{t-1})) \tilde{\chi}_{ijt}((\theta_{-i}^A, \theta_i^A), (\theta_{-i}^A, \theta_i^A), b_t, x^{t-1}) \right] \\ & \geq \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N_A \setminus \{i\}} \sum_{j \in N_B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}(x^{t-1})) \tilde{\chi}_{rjt}((\theta_{-i}^A, \hat{\theta}_i^A), (\theta_{-i}^A, \hat{\theta}_i^A), b_t, x^{t-1}) \right] \\ & + \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_i^A, \theta_{-i}^A} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} \left(\beta_i^A(\theta_i^A) b_{ijt}^A \frac{\theta_i^A}{\hat{\theta}_i^A} + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}(x^{t-1}) \right) \tilde{\chi}_{ijt}((\theta_{-i}^A, \hat{\theta}_i^A), (\theta_{-i}^A, \hat{\theta}_i^A), b_t, x^{t-1}) \right]. \end{aligned}$$

Combining the last two inequalities, and using the fact that horizontal and vertical types are independent, and that vertical types are drawn independently across agents, we have that

$$\left(\beta_i^A(\hat{\theta}_i^A) \hat{\theta}_i^A - \beta_i^A(\theta_i^A) \theta_i^A \right) \cdot \left(D_i^A((\theta_{-i}^A, \hat{\theta}_i^A); \tilde{\chi}) - D_i^A((\theta_{-i}^A, \theta_i^A); \tilde{\chi}) \right) \geq 0.$$

Because $\beta_i^A(\cdot)$ is strictly positive and non-decreasing, it must be that

$$D_i^A((\theta_{-i}^A, \hat{\theta}_i^A); \tilde{\chi}) \geq D_i^A((\theta_{-i}^A, \theta_i^A); \tilde{\chi}).$$

Part (ii). Since $\tilde{\chi}$ maximizes continuation weighted surplus after any history, arguments similar to those used to establish part (i) above imply that, for any $\hat{\theta}_{i0}^A, \theta_{i0}^A \in \Theta_i^A$, $\theta \in \Theta$,

$$\left(\beta_i^A(\hat{\theta}_{i0}^A) - \beta_i^A(\theta_{i0}^A) \right) \cdot \theta_{i0}^A \cdot \left(\tilde{D}_i^A(\hat{\theta}_{i0}^A, \theta; \tilde{\chi}) - \tilde{D}_i^A(\theta_{i0}^A, \theta; \tilde{\chi}) \right) \geq 0,$$

Because $\theta_{i0}^A > 0$ and $\beta_i^A(\cdot)$ is strictly positive and non-decreasing, it must be that $\tilde{D}_i^A(\hat{\theta}_{i0}^A, \theta; \tilde{\chi}) \geq \tilde{D}_i^A(\theta_{i0}^A, \theta; \tilde{\chi})$.⁴⁴

■

Next, for any agent $i \in N_A$ (the arguments for the side- B agents are analogous), let

$$\begin{aligned} \tilde{U}_i^A(\theta_{i0}^A; \theta) & \equiv \mathbb{E}^{\lambda[\chi]|\theta} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} \theta_i^A \varepsilon_{ijt}^A \chi_{ijt}((\theta_{-i}^A, \theta_{i0}^A), \theta_t, b_t, x^{t-1}) \right] \\ & - \mathbb{E}^{\lambda[\chi]|\theta} \left[\sum_{t=0}^{\infty} \delta^t \psi_{it}^A((\theta_{-i}^A, \theta_{i0}^A), \theta_t, b_t, x^{t-1}; \beta) \right] \end{aligned}$$

denote the payoff that the agent expects in the auctions defined by the rules (χ, ψ) , when the true vertical type profile is θ , the agent chooses the membership status θ_{i0}^A in period zero, he follows a truthful strategy from

⁴⁴Note that if $\beta_i^A(\hat{\theta}_{i0}^A) = \beta_i^A(\theta_{i0}^A)$, then $\tilde{D}_i^A(\hat{\theta}_{i0}^A, \theta; \tilde{\chi}) = \tilde{D}_i^A(\theta_{i0}^A, \theta; \tilde{\chi})$.

period $t = 1$ onwards, and all other agents follow truthful strategies from period $t = 0$ onwards.⁴⁵ Note that $\tilde{U}_i^A(\theta_i^A; (\theta_{-i}^A, \theta_i^A)) = \hat{U}_i^A(\theta_{-i}^A, \theta_i^A)$, where \hat{U} is as in (20).

From Step 2, under the rules $(\tilde{\chi}, \tilde{\psi})$, participating and following truthful strategies is a periodic ex-post continuation equilibrium starting from any period-1 history (including those reached off path, by deviations in period zero).⁴⁶ Standard arguments can then be used to show that the following envelope condition must be satisfied for all $l \in N_k$, $k = A, B$, all $\theta_{l_0}^k \in \Theta_l^k$, all $\theta \in \Theta$,

$$\tilde{U}_l^k(\theta_{l_0}^k; \theta) = \tilde{U}_l^k(\theta_{l_0}^k; (\theta_{-l}^k, \theta_{l_0}^k)) + \int_{\theta_{l_0}^k}^{\theta_l^k} \tilde{D}_l^k(\theta_{l_0}^k, (\theta_{-l}^k, y); \tilde{\chi}) dy. \quad (21)$$

The payoff that agent $l \in N_k$ from side $k = A, B$ obtains by selecting the membership $\theta_{l_0}^k$ when the true vertical type profile is θ is thus given by

$$\begin{aligned} \tilde{U}_l^k(\theta_{l_0}^k; \theta) &= \tilde{U}_l^k(\theta_{l_0}^k; (\theta_{-l}^k, \theta_{l_0}^k)) + \int_{\theta_{l_0}^k}^{\theta_l^k} \tilde{D}_l^k(\theta_{l_0}^k, (\theta_{-l}^k, y); \tilde{\chi}) dy \\ &\leq \tilde{U}_l^k(\theta_{l_0}^k; (\theta_{-l}^k, \theta_{l_0}^k)) + \int_{\theta_{l_0}^k}^{\theta_l^k} \tilde{D}_l^k(y, (\theta_{-l}^k, y); \tilde{\chi}) dy = \hat{U}_l^k(\theta_{-l}^k, \theta_{l_0}^k) + \int_{\theta_{l_0}^k}^{\theta_l^k} D_l^k((\theta_{-l}^k, y); \tilde{\chi}) dy \\ &= \hat{U}_l^k(\theta), \end{aligned}$$

where the first equality follows from (21), the inequality follows from part (ii) in Lemma 7, and the other equalities follow from (20) and the definition of the interim expected payoffs. Hence, given θ , the agent is better off following a truthful strategy from period zero onwards than deviating in period zero and then following a truthful strategy from period one onwards.

Finally, note that, since for any $l \in N_k$, $k = A, B$, any $\theta \in \Theta$, $\hat{U}_l^k(\theta)$ is bounded, participation constraints can always be satisfied by choosing the constants L_l^k appropriately.

Combining the results in Step 2 with those in Step 1, we thus have that, when the scoring rules and the associated payment functions satisfy the conditions in the theorem, participating and following a truthful strategy is a periodic ex-post equilibrium in the entire game. Q.E.D.

Proof of Theorem 2. Consider any feasible mechanism Γ and any BNE σ of the game induced by Γ . Denote by $\hat{\chi} = (\hat{\chi}_t(\theta, \varepsilon^t))_{t=1}^\infty$ and $\hat{\psi} = (\hat{\psi}_t(\theta, \varepsilon^t))_{t=1}^\infty$ the matching and payment rules, as a function of the true state, induced by σ in Γ . As mentioned in the main text, in case of endogenous processes, these functions are defined only for histories that are consistent with equilibrium play in previous periods. Also note that we are allowing here for *any* feasible mechanism; that is, the message and signal spaces may be different than those in the matching auctions.

⁴⁵The dependence of the payoff on the mechanism (χ, ψ) is omitted for convenience.

⁴⁶Recall that $\tilde{\chi}$ maximizes continuation weighted surplus and that the associated payment rule $\tilde{\psi}$ satisfies the conditions in (8) and (9) in the main text. From Step 1, $\tilde{\chi}$ coincides with either $\chi^{m;\beta}$, or with $\chi^{I;\beta}$, depending on whether the environment is the one in part (i) or part (ii) in the theorem.

The platform's profits under (Γ, σ) are equal to

$$\mathbb{E}^{\lambda[\hat{\chi}]} \left[\sum_{k=A,B} \sum_{l \in N_k} \sum_{t=0}^{\infty} \delta^t \hat{\psi}_{lt}^k(\theta, \varepsilon^t) - \sum_{t=1}^{\infty} \delta^t \sum_{i \in N_A} \sum_{j \in N_B} c_{ijt} (\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1})) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right], \quad (22)$$

where $\lambda[\hat{\chi}]$ denotes the process over vertical and horizontal types under the matching rule $\hat{\chi}$ induced by the strategies σ in Γ . Alternatively, (22) can be rewritten as follows:

$$\mathbb{E}^{\lambda[\hat{\chi}]} \left[\sum_{t=1}^{\infty} \sum_{i \in N_A} \sum_{j \in N_B} \delta^t ((\theta_i^A \varepsilon_{ijt}^A + \theta_j^B \varepsilon_{ijt}^B - c_{ijt} (\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1}))) \hat{\chi}_{ijt}(\theta, \varepsilon^t)) - \sum_{k=A,B} \sum_{l \in N_k} U_l^k(\theta_l^k) \right], \quad (23)$$

where $U_l^k(\theta_l^k)$ denotes the period-0 interim expected payoff of agent $l \in N_k$ from side $k = A, B$ when his vertical type is θ_l^k , under the equilibrium σ in the mechanism Γ . Note that we denote the interim payoffs by U_l^k to differentiate them from the interim payoff functions \hat{U}_l^k in the auction, whose domain is the entire profile of vertical types.⁴⁷

The period-0 participation constraints are satisfied if for all $l \in N_k$, $k = A, B$, $\theta_l^k \in \Theta_l^k$, $U_l^k(\theta_l^k) \geq 0$.

Following steps similar to those in Pavan, Segal and Toikka (2014, Theorem 1), we can show that the period-0 (interim) expected payoff of each agent $l \in N_k$, $k = A, B$ must satisfy the following envelope condition

$$U_l^k(\theta_l^k) = U_l^k(\underline{\theta}_l^k) + \int_{\underline{\theta}_l^k}^{\theta_l^k} \mathbb{E} \left[D_l^k(\theta; \hat{\chi}) | y \right] dy, \quad (24)$$

where the expectation is taken over the entire profile of vertical types θ given agent l 's own vertical type. This envelope condition, together with integration by parts, yields the following representation of the platform's profits,

$$\mathbb{E}^{\lambda[\hat{\chi}]} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i \in N_A} \sum_{j \in N_B} \left(\left(\theta_i^A - \frac{1 - F_i^A(\theta_i^A)}{f_i^A(\theta_i^A)} \right) \varepsilon_{ijt}^A + \left(\theta_j^B - \frac{1 - F_j^B(\theta_j^B)}{f_j^B(\theta_j^B)} \right) \varepsilon_{ijt}^B - c_{ijt} (\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1})) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right) \right] - \sum_{k=A,B} \sum_{l \in N_k} U_l^k(\underline{\theta}_l^k). \quad (25)$$

The first term, which is only a function of the matching rule $\hat{\chi}$, is the *expected dynamic virtual surplus* (DVS) generated by the matching rule $\hat{\chi}$. Clearly, because such a representation applies to the matching rule generated by any BNE of any mechanism, the above representation applies also to the state-contingent matching rules generated by the truthful strategies in the matching auctions.

Now observe that, when the weights are given by $\hat{\beta} \equiv (\hat{\beta}_l^k(\cdot))_{l \in N_k, k=A,B}$, the state-contingent matches implemented under the truthful equilibria of the auctions of Theorem 1 maximize DVS over all feasible state-contingent rules $\hat{\chi}$. More precisely, when the processes are exogenous, the matches implemented under a myopic scoring rule maximize DVS. Similarly, when the processes are endogenous and either $M = 1$, or $M \geq n_A \cdot n_B$,

⁴⁷ \hat{U}_l^k denotes the payoff an agent expects in the matching auction when the true vertical type profile is θ and all agents follow truthful strategies at all periods.

or the environment is separable, the matches implemented under the index rule maximize DVS. This is because (i) these matches have been shown to maximize the continuation weighted surplus (6) for any strictly positive and non-decreasing weights $\beta = (\beta_l^k(\cdot))_{l \in N_k, k=A, B}$ (step 1 in the proof of Theorem 1), (ii) the ex-ante weighted surplus when the weights are given by $\hat{\beta}$ coincides with DVS, and (iii) the weights $\hat{\beta}$ are strictly positive and non-decreasing.

Next, suppose the processes are exogenous. Let $\psi^{m; \hat{\beta}}$ be the payment scheme associated with the myopic scoring rule $\chi^{m; \hat{\beta}}$ with weights given by $\hat{\beta}$ (as defined by (8) and (9) in the main text) and in which $L_l^k = 0$, all $l \in N_k, k = A, B$. From (8) and (9), it is easy to see that, under the equilibria in truthful strategies of the matching auctions where (a) the scoring rule is the myopic one with weights given by $\hat{\beta}$, and (b) the payments are $\psi^{m; \hat{\beta}}$, the payoff expected by the lowest vertical type of each agent is exactly equal to zero (that is, $U_l^k(\underline{\theta}_l^k) = 0$, all $l \in N_k, k = A, B$). This means that the truthful equilibria of the above matching auctions maximize both terms of (25). Provided all the period-0 participation constraints are satisfied (something we verify below), we then have that the platform's profits are maximized under the truthful equilibria of the proposed matching auctions.

Also note that while we have restricted attention to deterministic mechanisms, the platform cannot increase its profits by using a randomized mechanism. This is because any randomized mechanism is equivalent to a deterministic one that conditions on the type realizations of a fictitious agent. The platform's profits under any equilibrium of such a mechanism thus continue to be given by the expression in (25), but with the matching rule conditioning on the type realizations of such fictitious agent. This means that the platform's profits are equal to the weighted average of the platform's profits under the deterministic matching rules obtained by conditioning on the various types of the fictitious agent. Because (25) is maximized over all possible deterministic rules under the equilibria in truthful strategies of the proposed matching auctions, we thus have that stochastic mechanisms can never improve upon the equilibria of the proposed auctions when it comes to the platform's profits.

We now complete the proof by establishing that all period-0 participation constraints are satisfied under the equilibria in truthful strategies of the proposed auctions. To see this, it suffices to observe that, for all $\theta_{-l}^k, l \in N_k, k = A, B$, $D_l^k((\theta_{-l}^k, \theta_l^k); m, \hat{\beta})$ is non-decreasing in θ_l^k , as established in Lemma 7 in the proof of Theorem 1. The assumption in the theorem that $D_l^k((\theta_{-l}^k, \underline{\theta}_l^k); m, \hat{\beta}) \geq 0$, all $l \in N_k, k = A, B$, all θ_{-l}^k then guarantees that $D_l^k(\theta; m, \hat{\beta}) \geq 0$, all $\theta \in \Theta, l \in N_k, k = A, B$. Because the period-0 interim payoffs satisfy the envelope condition

$$\hat{U}_l^k(\theta) = \int_{\underline{\theta}_l^k}^{\theta_l^k} D_l^k(\theta_{-l}^k, y; \hat{\chi}) dy + L_l^k.$$

(see the proof of Theorem 1), we then have that $\hat{U}_l^k(\theta) \geq 0$, all $\theta \in \Theta, l \in N_k, k = A, B$, which means that all the period-0 participation constraints are satisfied (in a periodic ex-post sense, i.e., for any θ , and not just in expectation over θ_{-l}^k given θ_l^k).⁴⁸

The optimality of the matching auctions with index rules in part (ii) of the theorem follow from similar arguments. Q.E.D.

⁴⁸Note that, under $\psi^{m; \hat{\beta}}$, $\hat{U}_l^k(\theta_{-l}^k, \theta_l^k) = 0$ all θ_{-l}^k .

Proof of Theorem 3. As explained in the main text, parts (i) and (ii) follow from Theorem 1. Part (iii) follows from arguments similar to those establishing the optimality of the auctions of Theorem 2. In particular, it follows from the fact that (a) the platform's expected profits under any BNE of any mechanism Γ implementing the welfare-maximizing matches satisfy the representation in (25); (b) each agent's period-0 expected payoff satisfies Condition (24); (c) in the proposed auctions, $U_l^k(\underline{\theta}_l^k) = 0$ if, and only if, the payments in (8) and (9) (for $\beta = \beta^W$) are such that $L_l^k = 0$, all $l \in N_k$, $k = A, B$, and (d) when the payments are given by (8) and (9) with $\beta = \beta^W$ and $L_l^k = 0$, all $l \in N_k$, $k = A, B$, all agents' period-0 participation constraints are satisfied, regardless of their beliefs over other agents' types, if, and only if, $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; m, \beta^W) \geq 0$, or $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; I, \beta^W) \geq 0$ (under exogenous or endogenous processes, respectively). The latter property is a result of the fact that the functions $D_l^k(\theta_{-l}^k, \cdot; m, \beta^W)$ and $D_l^k(\theta_{-l}^k, \cdot; I, \beta^W)$ are non-decreasing, which is shown in Lemma 7 in the proof of Theorem 1. Q.E.D.

Proof of Theorem 4. Let $\chi_t^P(\theta, \omega)$ and $\chi_t^W(\theta, \omega)$ denote the state-contingent matches implemented in period $t \geq 1$, under the truthful equilibria of, respectively, the profit-maximizing and the welfare-maximizing auctions of Theorems 2 and 3. Note that the arguments of these functions are the exogenous vertical types θ , and the sequences of exogenous innovations $\omega \equiv (\omega_{ijs}^k)_{(i,j) \in N_A \times N_B, k \in \{A, B\}}^{s=1, \dots, \infty}$ that, along with the matches implemented in previous periods (in case of endogenous processes), generate the horizontal types ε . As explained in the main text, such a representation favors the comparison of the matches sustained under the two auctions by making the "state" exogenous, thus eliminating the confusion that may originate from the fact that the histories of horizontal types need not coincide under the two auctions.

Similarly, let $S_{ijt}^{m;P}(\theta, \omega)$ and $S_{ijt}^{I;P}(\theta, \omega)$ denote, respectively, the period- t state-contingent myopic and index scores under the truthful equilibria of the profit-maximizing auctions of Theorem 2. Likewise, let $S_{ijt}^{m;W}(\theta, \omega)$ and $S_{ijt}^{I;W}(\theta, \omega)$ be the counterparts of $S_{ijt}^{m;P}(\theta, \omega)$ and $S_{ijt}^{I;P}(\theta, \omega)$ under the truthful equilibria of the welfare-maximizing auctions of Theorem 3. Because (θ, ω) are exogenous and time-invariant, they are dropped from all the functions $\chi^P, \chi^W, S^{m;P}, S^{m;W}, S^{I;P}$, and $S^{I;W}$ below.

First, observe that, because $\hat{\beta}_l^k(\theta_l^k) \leq 1 = \beta_l^{k,W}(\theta_l^k)$, all $\theta_l^k \in \Theta_l^k$, $l \in N_k, k = A, B$, and because the horizontal types are nonnegative, when processes are exogenous, for any $(i, j) \in N_A \times N_B$, $S_{ijt}^{m;P} \leq S_{ijt}^{m;W}$. Likewise, when processes are endogenous, for any $(i, j) \in N_A \times N_B$, any $t, \tau \geq 1$,

$$\sum_{s=1}^{t-1} \chi_{ijs}^W = \sum_{s=1}^{\tau-1} \chi_{ijs}^P \Rightarrow S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}. \quad (26)$$

Part 1. First, consider the case of exogenous processes. Because the capacity constraint is not binding, in each period $t \geq 1$, the matches implemented under the equilibria of the profit-maximizing auctions (alternatively, of the welfare-maximizing auctions) are all those for which the myopic scores $S_{ijt}^{m;P} \geq 0$ (alternatively, $S_{ijt}^{m;W} \geq 0$). The above property, along with the fact that, for any $(i, j) \in N_A \times N_B$, $t \geq 1$, $S_{ijt}^{m;W} \geq S_{ijt}^{m;P}$, then yields the result.

Next, consider the case of endogenous processes. Again, because the capacity constraint is not binding, in each period $t \geq 1$, the matches implemented under the equilibria of the profit-maximizing auctions (alternatively, the welfare-maximizing auctions) are all those for which the index $S_{ijt}^{I;P} \geq 0$ (alternatively, $S_{ijt}^{I;W} \geq 0$). The result then follows from the fact that, for any $(i, j) \in N_A \times N_B$, $t \geq 1$, $S_{ijt}^{m;W} \geq S_{ijt}^{m;P}$. The last property,

in turn, follows by induction. First observe that the property is necessarily true at $t = 1$, given (26) and the fact that, in period $t = 1$, the number of past interactions is necessarily the same under profit and welfare maximization. Now suppose the result holds for all $1 \leq s < t$. Note that any match for which $S_{ijt}^{I;P} \geq 0$ has been active at each preceding period $s < t$, both under profit maximization and under welfare maximization. The result then follows again from (26), which implies that $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$.

Part 2. Consider first the case of exogenous processes and $M < n_A \cdot n_B$. The result follows directly from the following two properties: (a) in each period $t \geq 1$, the set of matches for which $S_{ijt}^{m;P} \geq 0$ is a subset of the set of matches for which $S_{ijt}^{m;W} \geq 0$, (b) the cardinality of the set of matches implemented in each period in a profit-maximizing auction (alternatively, in a welfare-maximizing auction) is the minimum between M and the cardinality of the set of matches for which $S_{ijt}^{m;P} \geq 0$ (alternatively, $S_{ijt}^{m;W} \geq 0$).

Next, consider the case of endogenous processes with $M = 1$. First observe that, under the equilibria of the profit-maximizing auction, if at some period $t \geq 1$, $\chi_{ijt}^P = 0$, all $(i, j) \in N_A \times N_B$, then $\chi_{ijs}^P = 0$, all $s > t$, all $(i, j) \in N_A \times N_B$. The same property holds for χ^W . Next, observe that, if matching stops at period t under profit maximization (alternatively, welfare maximization), then $S_{ijt}^{I;P} < 0$ all $(i, j) \in N_A \times N_B$ (alternatively, $S_{ijt}^{I;W} < 0$ all $(i, j) \in N_A \times N_B$). Now suppose that, under profit maximization, matching is still active in period t (meaning, there exists $(i, j) \in N_A \times N_B$ such that $\chi_{ijt}^P = 1$). Then there are two cases. (1) Either $\sum_{s=1}^{t-1} \chi_{ijs}^W = \sum_{s=1}^{t-1} \chi_{ijs}^P$, all $(i, j) \in N_A \times N_B$, in which case (26) implies that $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$ for all $(i, j) \in N_A \times N_B$, which implies the result. Or, (2) there exists $(i, j) \in N_A \times N_B$ such that $\sum_{s=1}^{t-1} \chi_{ijs}^W < \sum_{s=1}^{t-1} \chi_{ijs}^P$. In this case, there must exist $\tau < t$ such that $\sum_{s=1}^{\tau-1} \chi_{ijs}^P = \sum_{s=1}^{\tau-1} \chi_{ijs}^W$, and $\chi_{ij\tau}^P = 1$. That $\chi_{ij\tau}^P = 1$ in turn implies $S_{ij\tau}^{I;P} \geq 0$. That $\sum_{s=1}^{\tau-1} \chi_{ijs}^P = \sum_{s=1}^{\tau-1} \chi_{ijs}^W$ in turn implies that $S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}$, where the result follows again from (26). From the discussion above, that $S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}$ in turn implies that matching must be active in period t also under welfare maximization.

Part 3. Since, under profit maximization, matching terminates after a finite number of periods $T \in \mathbb{N}$, at period $T + 1$, the index $S_{ijT+1}^{I;P} < 0$ all $(i, j) \in N_A \times N_B$. Take any pair (i, j) and denote the number of times the match (i, j) has been active under profit maximization prior to period $T + 1$ by R_{ij}^P . For any $n \leq R_{ij}^P$, the index $S_{ij}^{I;P}$ at the n -th time the pair was matched must have been nonnegative. The same must then be true for the index $S_{ij}^{I;W}$ (this follows again from (26)). Therefore, matching will not stop under welfare maximization until the pair (i, j) is matched at least R_{ij}^P times. Since this holds for each pair $(i, j) \in N_A \times N_B$, the result in Part (3) follows. Q.E.D.

Details for the example in Section 2. We first show that the indexes in the example correspond to the indexes in (5). Observe that, in each period $t \geq 1$, beliefs evolve according to the following process. If $(i, j) \in N_A \times N_B$ were not matched in period t , then $\eta_{ijt}^k = \eta_{ijt-1}^k$, $k = A, B$. If, instead, (i, j) were matched at $t - 1$ and $i \in N_A$ (similarly, $j \in N_B$) received a negative signal in any of the previous periods, then $\eta_{ijt}^A = 0$. Finally, if (i, j) were matched at $t - 1$, and in any of the previous periods i received positive signals, then

$$\eta_{ijt}^A = \frac{\eta_{ijt-1}^A}{\eta_{ijt-1}^A + (1 - \eta_{ijt-1}^A)\mu_{ij}} = \frac{\eta_{ijt-1}^A}{\eta_{ijt-1}^A(1 - \mu_{ij}) + \mu_{ij}}.$$

By induction, given m previous interactions with $j \in N_B$ in which $i \in N_A$ did not receive a negative signal, i 's

period- t posterior belief is equal to

$$\eta_{ijt}^A(m) = \frac{\eta_{ij0}}{\eta_{ij0} \left(1 - \mu_{ij}^m\right) + \mu_{ij}^m}.$$

Similarly, i 's flow payoff given m past interactions with j after having received only positive signals is equal to

$$v_{ijt}^A(m) = (2\eta_{ijt}^A(m) - 1) u_{ij} = \left(\frac{\eta_{ij0} - \mu_{ij}^m (1 - \eta_{ij0})}{\eta_{ij0} \left(1 - \mu_{ij}^m\right) + \mu_{ij}^m} \right) u_{ij}.$$

At any period t , given $m - 1$ past interactions with j , the probability that i , after having observed only positive signals in the past, receives an additional positive signal in period t after being matched once more with j (for the m 'th time) is equal to

$$\kappa_{ijt}(m) \equiv \eta_{ijt}^A(m) + \mu_{ij} \left(1 - \eta_{ijt}^A(m)\right) = \frac{\eta_{ij0} \left(1 - \mu_{ij}^m\right) + \mu_{ij}^m}{\eta_{ij0} \left(1 - \mu_{ij}^{m-1}\right) + \mu_{ij}^{m-1}}.$$

We can now compute the index (5). Suppose that, after having interacted m times in the past, both agents i and j submit a positive bid. We then have, after some algebra,⁴⁹

$$\begin{aligned} S_{ijt} &= 2 \frac{v_{ijt}^k(m) + \sum_{s=1}^{\infty} \left(\delta^s v_{ijt}^k(m+s) \prod_{r=1}^s (\kappa_{ijt}(m+r))^2 \right)}{1 + \sum_{s=1}^{\infty} \left(\delta^s \prod_{r=1}^s (\kappa_{ijt}(m+r))^2 \right)} - c_{ij} \\ &= 2u_{ij} \left(\frac{\frac{\eta_{ij0}^2}{1-\delta} - \frac{(1-\eta_{ij0})^2 \mu_{ij}^{2m}}{1-\delta \mu_{ij}^2}}{\frac{\eta_{ij0}^2}{1-\delta} + \frac{2\eta_{ij0}(1-\eta_{ij0})\mu_{ij}^m}{1-\delta \mu_{ij}} + \frac{(1-\eta_{ij0})^2 \mu_{ij}^{2m}}{1-\delta \mu_{ij}^2}} \right) - c_{ij}. \end{aligned}$$

Similarly, if one of the two agents submits a negative bid, then

$$\begin{aligned} S_{ijt} &= \frac{v_{ijt}^k(m) + \sum_{s=1}^{\infty} \left(\delta^s v_{ijt}^k(m+s) \prod_{r=1}^s \kappa_{ijt}(m+r) \right)}{1 + \sum_{s=1}^{\infty} \left(\delta^s \prod_{r=1}^s \kappa_{ijt}(m+r) \right)} - u_{ij} - c_{ij} \\ &= -2u_{ij} \left(\frac{\frac{\mu_{ij}^m (1-\eta_{ij0})}{1-\delta \mu_{ij}}}{\frac{\eta_{ij0}}{1-\delta} + \frac{\mu_{ij}^m (1-\eta_{ij0})}{1-\delta \mu_{ij}}} \right) - c_{ij}. \end{aligned}$$

The latter indexes correspond to those in the example.

Given the assumptions in the example, in period $t = 1$: $S_{111} \approx 0.299$, $S_{121} \approx 1.895$ and $S_{211} \approx 1.885$. The pairs (1,2) and (2,1) are therefore matched at $t = 1$. Furthermore, because $\delta = 4/5$, separability is clearly satisfied at $t = 1$. Each of these matches is therefore implemented until a negative bid is submitted by one of the involved agents, at which point the corresponding match is replaced with (1,1). When a negative bid is submitted in either of the matches (1,1) or (2,1), that match is never implemented again, as its index will remain negative forever. The index corresponding to (1,2) remains positive after a single negative bid is received as long as the number of past matches between agent 1 from side 1 and agent 2 from side 2 before one of the two

⁴⁹The assumption that $\eta_{ij0} \geq \frac{1}{2}$ is sufficient to guarantee that the stopping time is the correct one.

agents receives a negative signal is greater than $m = 2$. It remains to verify that separability is satisfied at all periods $t > 1$. In the environment considered in the example, we need only to check that the *initial* index corresponding to the match (1,1), multiplied by $(1 - \delta)$, is greater than the index for the match (1,2), given any number of past interactions m between agent 1 from side 1 and agent 2 from side 2, which is indeed the case. Q.E.D.

References

- [1] Akbarpour, M., Li, S., and S. Oveis Gharan (2016), “Thickness and Information in Dynamic Matching Markets,” mimeo Stanford University.
- [2] Allard, N., (2008), “Lobbying Is An Honorable Profession: The Right To Petition and The Competition To Be Right,” *Stanford Law Review*, Vol. 19(1), pp. 23-68.
- [3] Anderson, R., Ashlagi, I., Gamarnik, D., and Y. Kanoria (2015), “Efficient Dynamic Barter Exchange,” mimeo, MIT.
- [4] Asker, J., and E. Cantillon (2008), “Properties of Scoring Auctions,” *The RAND Journal of Economics*, 39, pp. 69-85.
- [5] Athey, S. and G. Ellison (2011), “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, Vol. 126, pp. 1213–1270.
- [6] Athey, S. and I. Segal (2013), “An Efficient Dynamic Mechanism,” *Econometrica*, 81(6), 2463–2485.
- [7] Baccara, M., S. Lee, and L. Yariv (2015): “Optimal Dynamic Matching,” mimeo Caltech.
- [8] Baron, D. and D. Besanko (1984), “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1, 447–470.
- [9] Bedre-Defolie, O. and E. Calvano, (2013), “Pricing Payment Cards,” *American Economic Journal: Microeconomics*, Vol. 5(3), pp. 206-231.
- [10] Bergemann, D. and A. Pavan (2015), “Introduction to Symposium on Dynamic Contracts and Mechanism Design,” *Journal of Economic Theory* 159, 679-701.
- [11] Bergemann, D. and M. Said (2011), “Dynamic Auctions,” *Wiley Encyclopedia of Operations Research and Management Science*.
- [12] Bergemann, D. and P. Strack (2015), “Dynamic Revenue Maximization: A Continuous Time Approach,” *Journal of Economic Theory* 159, 819-853.
- [13] Bergemann, D. and J. Välimäki (2001), “Stationary Multi-Choice Bandit Problems,” *Journal of Economic dynamics and Control*, 25(10), 1585-1594.

- [14] Bergemann, D. and J. Välimäki (2008), “Bandit Problems,” *New Palgrave Dictionary of Economics*, Durlauf and Blume eds.
- [15] Bergemann, D. and J. Välimäki (2010), “The Dynamic Pivot Mechanism,” *Econometrica*, 78(2), 771–789.
- [16] Bergemann, D. and J. Välimäki (2017), “Dynamic Mechanism Design: An Introduction,” mimeo, Yale University.
- [17] Besanko, D. (1985), “Multiperiod Contracts Between Principal and Agent with Adverse Selection,” *Economics Letters*, 17, 33–37.
- [18] Bloch, F. and D. Cantala (2016), “Dynamic Allocation of Objects to Queuing Agents,” *American Economic Journal: Microeconomics*, forthcoming.
- [19] Board, S. (2007), “Selling Options,” *Journal of Economic Theory* 136, 324–340.
- [20] Board, S. (2009), “Monopoly Group Design with Peer Effects,” *Theoretical Economics*, Vol. 4, 89-125.
- [21] Bonatti, A. and J. Horner (2015), “Learning to Disagree in a Game of Experimentation,” mimeo Yale University.
- [22] Borgers, T. (2015), “An Introduction to the Theory of Mechanism Design,” Oxford University Press, Oxford.
- [23] Borgers, T., Cox, I., Pesendorfer, M., and V. Petricek (2013), “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence,” *American Economic Journal: Microeconomics*, Vol. 5, pp. 163-187.
- [24] Cabral L. (2011): “Dynamic Price Competition with Network Effects,” *Review of Economic Studies* 78 (1): 83-111.
- [25] Che, Y. K. (1993): “Design Competition through Multi-Dimensional Auctions,” *RAND Journal of Economics* Vol. 24, pp. 668-680.
- [26] Cremer, J. and R. P. McLean (1988), “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56(6), 1247–1257.
- [27] Courty, P. and H. Li (2000), “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- [28] Damiano, E. and H. Li (2007), “Price Discrimination and Efficient Matching,” *Economic Theory*, Vol. 30, pp. 243-263.
- [29] Damiano, E. and R. Lam, (2005), “Stability in dynamic matching markets,” *Games and Economic Behavior*, 52(1):34–53
- [30] Dekel, E., Jackson, M.O. and A. Wolinsky, (2008) “Vote Buying: General Elections,” *Journal of Political Economy*, 116(2), pp. 351–380.

- [31] Doepke, M and R.M. Townsend, (2006) “Dynamic Mechanism Design with Hidden Income and Hidden Actions,” *Journal of Economic Theory* 126, 235 – 285.
- [32] Doval, L. (2016), “A Theory of Stability in Dynamic Matching Markets,” mimeo Northwestern University.
- [33] Edelman, B., Ostrovsky, M., and M. Schwarz (2007), “Internet Advertising and the Generalized Second-Price Auction,” *American Economic Review*, Vol. 97, pp. 242-259.
- [34] Eső, P. and B. Szentes (2007), “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74(3), 705–731.
- [35] Fershtman, D. and A. Pavan (2017), “Pandora’s Auctions: Dynamic Matching with Unknown Preferences,” *American Economic Review Papers and Proceedings*.
- [36] Ishikida, T., and P. Varaiya (1994), “Multi-Armed Bandit Problem Revisited,” *Journal of Optimization Theory and Applications*, 83(1), 113-154.
- [37] Gershkov, A. and B. Moldovanu (2014), “Dynamic Allocation and Pricing: A Mechanism Design Approach,” MIT Press.
- [38] Gittins, J. C. (1979), “Bandit processes and dynamic allocation indices,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–177.
- [39] Gomes, R. (2014), “Optimal auction design in two-sided markets,” *The RAND Journal of Economics*, 45: 248–272.
- [40] Gomes, R. and A. Pavan (2016), “Many-to-Many Matching and Price Discrimination,” *Theoretical Economics*, 11(3), 1005–1052.
- [41] Gomes, R. and A. Pavan (2017), “Price Customization and Targeting in Many-to-Many Matching Markets,” mimeo Northwestern University and the Toulouse School of Economics.
- [42] Gomes, R. and A. Sweeney (2014), “Bayes–Nash Equilibria of the Generalized Second-price Auction,” *Games and Economic Behavior*, 86, 421–437.
- [43] Herbst, H., and B. Schickner (2016), “Dynamic Formation of Teams: When Does Waiting for Good Matches Pay Off?,” mimeo University of Bonn.
- [44] Horner, J., and A. Skrzypacz (2016), “Learning, Experimentation and Information Design,” mimeo Yale University.
- [45] Johnson, T., (2013), “Matching Through Position Auctions,” *Journal of Economic Theory*, 148, 1700-1713.
- [46] Jullien, B. and A. Pavan (2017), “Information Management and Pricing in Platform Markets,” mimeo Northwestern University and Toulouse School of Economics.

- [47] Kakade, S., Lobel, I., and H. Nazerzadeh (2011), “Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism,” *Operations Research*, 61, 837–854.
- [48] Kang, K. and H. You (2016), “Lobbyists as Matchmakers in the Market for Access,” mimeo Carnegie Mellon University.
- [49] Kapicka, M. (2013), “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach,” *Review of Economic Studies*, 80(3), 1027-1054.
- [50] Keller, G. and S. Rady (2015), “Breakdowns,” *Theoretical Economics*, 10(1), 175-202.
- [51] Keller, G., Rady, S. and M. Cripps (2005), “Strategic experimentation with exponential bandits,” *Econometrica*, 73(1), 39-68.
- [52] Kurino, M. (2009) “Credibility, efficiency and stability: a theory of dynamic matching markets,” *Jena economic research papers, JENA*.
- [53] Lee, R. (2013), “Vertical Integration and Exclusivity in Platform and Two-Sided Markets,” *American Economic Review*, Vol. 103(7), pp. 2960-3000.
- [54] Leshno, J. (2015), “Dynamic Matching in Overloaded Waiting Lists,” mimeo Columbia University.
- [55] Li, S., M. Mahdian, and P. McAfee (2010) “Value of Learning in Sponsored Search Auctions,” *Proceedings of the 6th international Workshop on Internet and Network Economics (WINE)*.
- [56] Makris, M., and A. Pavan (2017), “Taxation under Learning by Doing,” mimeo Northwestern University.
- [57] Mandelbaum, A. (1986), “Discrete Multi-Armed Bandits and Multi-Parameter Processes,” *Probability Theory and Related Fields*, 71(1), 129-147.
- [58] Mansour, Y., S. Muthukrishnan, and N. Nisan (2012) “Doubleclick Ad Exchange Auction,” working paper, Tel Aviv, Rutgers, and Hebrew University.
- [59] Moghaddam, M , and S. Y. Nof (2016), “Best Matching: Theory and Applications,” Springer.
- [60] Pandelis, D. G., and D. Teneketzis (1999), “On the Optimality of the Gittins Index Rule for Multi-Armed Bandits with Multiple Plays,” *Mathematical Methods of Operations Research*, 50(3), 449-461.
- [61] Pavan, A. (2017), “Dynamic Mechanism Design: Robustness and Endogenous Types,” *Advances in Economics and Econometrics*, Proceedings of the 11th World Congress of Econometric Society, Volume 1, 1-62.
- [62] Pavan, A., I. Segal, and J. Toikka (2014), “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.
- [63] Rysman, M., (2009), “The Economics of Two-Sided Markets,” *Journal of Economic Perspectives*, Vol. 23(3), pp. 125-143.

- [64] Schummer, J., (2016), “Influencing Waiting Lists,” mimeo Northwestern University.
- [65] Thakral, N., (2015), “Matching with Stochastic Arrival: Theory and an Application to Public Housing,” mimeo Harvard University.
- [66] Ünver, U. (2010), “Dynamic Kidney Exchange,” *Review of Economic Studies*, 77, 372- 414.
- [67] Varian, H. (2007), “Position Auctions,” *International Journal of Industrial Organization*, Vol. 25, pp. 1163-1178.
- [68] Weber, R. (1992). “On the Gittins index for multiarmed bandits,” *The Annals of Applied Probability*, 2(4), 1024-1033.
- [69] Weber, R. R., and G. Weiss (1990), “On an Index Policy for Restless Bandits,” *Journal of Applied Probability*, 637-648.
- [70] Weyl, G. (2010), “A Price Theory of Two-Sided Markets,” *American Economic Review*, Vol. 100(4), pp. 1642-72.
- [71] Whittle, P. (1982), “Optimization Over Time,” John Wiley & Sons, Inc.
- [72] Whittle, P. (1988), “Restless Bandits: Activity Allocation in a Changing World,” *Journal of applied probability*, 287-298.